Opening Presentation

COMPUTATIONAL SOCIAL SCIENCE: AGENTS, INTERACTION, AND DYNAMICS

K.M. CARLEY, Carnegie Mellon University*

ABSTRACT

Computational techniques are an important tool in the social theorists' and methodologists' toolkit. In particular, multi-agent computational models are particularly valuable for the development of social theory. They enable the researcher to examine the relations among groups and individual agents and to witness the way in which these relations enable, constrain, and affect agent and group behavior, as well as emergent phenomena. During the past 40 years, computational analysis more generally has played a role in revolutionizing social theory. Important advances in the area of cognition, interaction, chance, and adaptation have been made, leading to new paradigms such as information processing theory. Fundamental results related to bounded rationality, satisficing, competency traps, emergent order, and learning clashes have emerged. Today, multi-agent models are enabling social scientists to ask fundamental questions about the nature of coordination, mechanisms for facilitating or inhibiting change, and the effect of scale and technology on social behavior. Results and findings are illustrated using a variety of classical and current models.

[edited transcript of presentation follows]

INTRODUCTION

A couple of years ago I had the opportunity to go to a naval war game. Now, that may not sound like a really cool thing to most of you, but it was actually quite interesting. But the first thing that happened was that my portable computer got knocked off a desk onto the concrete and exploded all over the place. The military guy who was in charge turned to me and said, "Thank God. One less agent I have to get a security clearance for."

You may think that's a dumb remark, but it was true, in a sense. We were looking at new technologies in the military, and we had a lot of agent-based models. And a lot of them, of course, were stored on my machine. So it really did have to go through not a security clearance, exactly, but a virus check. One might begin to wonder about the extent to which agent models truly are social agents. That's the underlying theme of what I'm going to be talking about.

It has become increasingly clear that computational techniques in general are an extremely useful tool for doing social theorizing and for doing methodological development

* *Corresponding author address*: Kathleen M. Carley, Carnegie Mellon University, Porter Hall 219a, Pittsburgh, PA 15213; e-mail: kathleen.carley@cmu.edu.

[*slide 2*].[1] In particular, the new multi-agent techniques — and don't ask me what an agent is — are particularly valuable. They enable the researcher to do a lot of things. If you read any paper in this area, you'll see people talking about how agents allow you to examine relations among groups and individuals, let you witness the way relations enable or constrain individual behavior, and, of course, let you do emergent phenomena most readily.

Some of the key areas where they've been particularly useful in the social sciences are studies of the fundamental nature of coordination and communication, facilitating or inhibiting change among agents, the effect of scale (like moving from worlds of five to worlds of twenty to worlds of thousands), and the effect, of course, of information technology or telecommunications technology.

Now, over the past 40 years of computer models in this area, we've seen important advances in areas such as the cognitive sciences, interaction, chance, chaos, and adaptation [*slide 3*]. Fundamental results out there spring from ideas as simple as the old notion of bounded rationality, as well as notions of satisficing, competency traps, emergent order, learning clashes, and so on. The computational techniques have been responsible in large measure for the evolution of new paradigms, including the information processing paradigm and now, of course, the neo-information processing paradigm. The reason for this influence, many people will argue, is that the techniques allow you to ask new questions and to explore fundamental assumptions and to break the assumptions and try new ones. Next I'm going to go through some of the basic findings.

NATURE OF MODELS

The first thing I thought it might be instructive to look at is a set of models that are out there in the workplace [*slide 4*]. Clearly, this isn't all the models that exist in the field. It's not even close. But these models have appeared in multiple papers, and each of them has some claim to fame. Most are from the organizations area. What I've done here is take each of the models and show the extent to which they model various characteristics, like learning, multi-agency, features of tasks, features of organizational structure, resources, and so on.

Two things should be readily apparent in this table. First, a lot of attention has been given to multi-agent-type models. Many of these models have multiple agents in some form. The second thing, which is even more readily apparent, is that in addition to paying attention to individual agents, they've paid attention to resources, that is, the kinds of things the agents can manipulate, play with, change, adapt, alter, etc. They've paid a fair amount of attention to knowledge, but not as much, as you can see by the fact that not many deal with learning.

These models, plus the countless models I haven't put up there, have been used in a variety of ways. The list I've got here is not exhaustive [*slide 5*]. Nor are these strict alternatives in a logical sense. But these are some of the modes in which people are using computational models and have used them in the past 40 years.

---

[1] Slide numbers refer to the presentation slides that are reproduced beginning on page 15. Slides 10, 16, 20, and 23-28 are not referenced, as they have been omitted from this publication.

Models as theory. The first mode is probably the most profound and the most irritating to noncomputational people, and that's the notion of the model as a theory. The argument dates back to early computer science. It holds that the model does the task it seeks to explain; therefore, it is a theory. As such, one needs to test it in new ways. An example coming out of that approach is the social Turing test that Newell and I worked on, which extends the Turing tests to a set of agents and says, "Well, if you're using the simulation model in an experiment and it generates socially realistic behavior, then it meets the social Turing test." There's a little more to it, but that's about it.

Models as agents. In the second mode, models are basically used as agents. Ever since the very beginning of computational modeling techniques in the social sciences, we've had agent-like models. We've had models running around doing things, like playing in tournaments. We've had them taking part in experiments, substituting for humans in the lab, and so on. Not always were these terribly sophisticated, but models have played the role of individual agents in individual ways.

Models as virtual worlds. Many models create virtual worlds. The most familiar example to a lot of people is the A-Life work, particularly the work on Sugarscape at Brookings. The idea here is that you're doing social theorizing from the ground up: you're building worlds out of tons of these little agents interacting, talking to each other, communicating, and building emerging social phenomena out of their interactions.

Models as empirically grounded theories. Another mode follows a tradition that's extremely strong within the social sciences compared to other areas where computational techniques are used. In this mode, one builds empirical data into the simulation model, not only as parameters, but also as input data and so on, right as an integral part of the model. An example of that is the work by David Heise and his group on affect control theory and their model of social exchange as a result of affective changes in personality. Built into the model are the massive data sets collected from multiple cultures from multiple countries on the way individuals respond in various emotional settings, the emotional rating of various words, roles, positions, and situations, etc. They use that data to predict how these things will behave over time.

The thing I love about this approach is that a lot of the data embedded in simulation models today is ethnographic in nature. That is, if I were going to build model of a particular process at GM, the first thing I'd want is a very detailed ethnographic study of that group, because that would give me the details I need to build a simulation model.

Models as hypothesis generators. The models that you have nowadays are often so large that their responsiveness — that is, their performance outcomes relative to all the different kinds of input data — is just really large, and you can't always analyze the whole thing. So you run virtual experiments — like a human laboratory experiment, but with a computer — and statistically analyze the results. The end result is a set of hypotheses that can then be tested in other settings or compared against other models.

One of the values of that strategy is that it allows you to generate hypotheses from your basic theory in a way that is more systematic than verbal theorizing. In fact, when some researchers have translated existing verbal theories into simulations, they've found that the predictions and the hypotheses coming out of the verbal theories were actually incompatible with the fundamental assumption. It could not be regenerated. So it's a useful technique.

COMPUTATIONAL ANALYSIS

Multi-agent models. The computational approach a lot of people are using is to build multi-agent models. In these models, the agents can be models of humans or of artificial agents, like web-bots or robots or avatars or databases, and these agents tend to be diverse [*slide 6*]. That is, you may have multiple types of agents in a model, or they may be diverse simply by virtue of having access to different knowledge or, in the case of Sugarscape, sugar, or whatever. So you've got agent heterogeneity, which is really important, because a lot of the outcomes that are derived are really the result of having a set of heterogeneous rather than homogeneous agents.

Socio-information processing models. Social action is often derived through interaction among the agents, not by individual outcomes. And a lot of the models — in fact, I would say the vast bulk of them, particularly in organizations — are socio-information processing models. That is, the agents are information processors with the ability to collect, analyze, distribute, communicate, or do whatever they want with this information, but they're also "socio" in the sense that they usually have a little bit of a model of the other [agents].

Because they're socio-information processing agents, they have a set of constraints that really dictate behavior. So you're really talking about constrained behavior. The oldest constraint, of course, is cognitive constraints on the way you think, analyze, process data, etc., but the other ones are, of course, social constraints (who you talk to, who's in your network) and technological constraints (what technology you have access to for communication, for travel, for moving).

Emergent behavior. Out of these models you get emergent behavior. You can get it at the individual, group, organization, or population level or at multiple levels at once. You can also talk about the co-evolution of those systems. Emergent behavior is often studied by detecting patterns, and one of the difficulties in this area is doing pattern detection. That is, the ideas that people have and the models they're generating at this point are outstripping our ability methodologically to extract patterns, especially patterns that are dynamic over time, particularly for large-scale groups.

Equilibrium. Another feature of these models is that you rarely run them at equilibrium. You're usually not interested in questions of equilibrium; you're concerned with questions of dynamics and of change. In fact, a lot of the modelers right now will argue there's no such thing as an equilibrium; that's just a false assumption. These models also often tend to be empirically grounded at some level. At least they're based on empirical findings. The parameters may be input from other data sets, and so on. One of the implications is that you need multi-level validation and multiple types of validation.

CRITICAL FINDINGS IN FORTY YEARS OF RESEARCH

With that as background on the field in general, what I want to go through now is a series of critical findings from the past 40 years in this area. I think that over this time period a lot of really nice and really important results have come out, but they've come out of this department over here and that discipline over there. But if we put them all together, they're really suggesting a whole new way of thinking about human action and social behavior [*slide 7*]. So computational analysis has generated this entirely new way of thinking about it. The work that's been done in this area is strongly interdisciplinary. Many of the critical findings are the result of teams of people working together, anywhere from four to 30 people in a team, and, as I said, from many disciplines.

Cognitive Constraints

One of the oldest findings is that, in terms of cognition, there are constraints, and they really do matter [*slide 8*]. Now, in the old paradigm — because we're talking about the 1950s, basically — the view of the world is that you have rational actors. But look at one of the first models, the "garbage can" model of Cohen, March, and Olson. They said, "Look, these actors aren't rational." Human beings are not that way. They're boundedly rational at best. They have limits on their ability to process information, and moreover, what they process is in part affected by what they're interested in. A similar argument is made in the behavioral theory of the firm, which is a much more detailed, elaborate model by Cyert and March. And in those two models alone, they were able to demonstrate quite conclusively that if you use boundedly rational actors, the fundamental results that you might get from assuming rational actors just go by the wayside, and behavior looks totally different.

So, for example, the results might show that in organizations most decisions, rather than being made by finding the right decision, are actually made by oversight, by accident, because you happen to be in the right place at the right time — and that data actually reflects the world. In other words, they showed that cognitive and social limitations really affect the social outcomes.

Architecture

The second finding, which is a further refinement of the first, was the idea that in terms of cognition, the actual architecture matters [*slide 9*]. The term "architecture" from cognitive science refers to the actual hard wiring in your brain for how you process and handle information. If you're talking about a computer program, it's the actual code for how it handles information.

In terms of architecture, it's not that there was one paradigm; there were a whole bunch of views out there. There's the strong structural position, coming out of sociology: you have networks; they really matter; if you replace the human beings with rocks or rats, it doesn't matter — you'll get the same effect. Then there's the rational actor view that says agent differences just don't matter; they're not relevant. And there's the decision-making view that says outcomes are based on decision processes and the information gathering has no impact. But in fact the simulation work in this area said, "Well, that's not true."

What our models are suggesting is that agent variation exists, and the variation comes at two levels. First, it comes in terms of information processing capabilities. That is, if I build a model with certain capabilities, with a certain way of handling information, and then I build another one, they may generate fundamentally different results. Experiments with humans and rats show that the same situation will generate fundamentally different results because the cognitive architecture matters. The second reason why it matters is that agents are heterogeneous not just in terms of their cognitive architecture but also in terms of the knowledge that they have, and differences in knowledge also matter.

Technology

The findings about cognitive architecture opened a door to another finding, which is that we can treat any kind of telecommunication technology. Any information processing technology itself adds an agent with information processing capabilities that are distinct from those of humans. When you put humans and these artificial agents together, you will get distinctive differences in the outcomes because the exact way they are processing the information is different. Here again, it's not just that constraints matter, as we just saw, but now it's that the *exact* constraints totally affect the outcomes.

Another way of looking at this is to say that agents don't have to be people. Here's an example. We're looking at information diffusion, namely, the time for information to diffuse relative to how professional a group is. The group is more professional if it has more information than the average group or average society *and* if that information is select and special to that group. So there is a set of knowledge that they know and that is more or less exclusive to their group.

A good example of a professional society would be physics. [At first] you have a world with just human beings who can communicate information. [Then you look at] what happens when you add additional agents, which are their web pages, and how that affects the diffusion of information. [We see] that it takes the longest for information to diffuse (the highest point) when the new idea originates in a very professional group and it's going to a very unprofessional group. Think of this as physics going to the general public; it takes a long time. But when you introduce web pages, the whole surface changes. And now where did information move the fastest? Physics to the general public. Some people say this is a good explanation for what happened with the cold fusion story.

Interactions

So we've talked about cognition. Another huge area where there has been a lot of research is interaction [*slide 11*]. The earliest work in this area was associated with games. In particular, it was associated with the Prisoner's Dilemma game. The basic issue was the type of social intelligence needed to get groups to cooperate. You all know of Axelrod's tournaments, where people came in and played the Prisoner's Dilemma, and they sent in a computer program. Much to his and everyone else's surprise, a simple tit-for-tat strategy won in the first tournament.

That was very important because it showed the value of the reciprocity norm. That result has since inspired a lot of research, including work showing that interesting social orders, like cooperation, can emerge even with zero-intelligence actors; that the order in which the individuals interact matters to the outcome; that you can generate different behavior if you allow people to alter the choices.

This one particular finding, that interaction matters, led to a huge range of other findings in this area of interaction. But the fundamental idea is that through simulation, through these agent models, we were able to see, in fact, that interaction becomes a critical determinant of social outcome. So it is not just cognition; it is interaction, too.

Networks

More recently, people have been working on networks. The argument is that networks matter [*slide 12*]. I'm sure you have all heard about the "small world" phenomenon. The basic idea there is that we have a set of people who can interact. Think of the old rumor game. They are communicating information, and let's say the red guy there thinks up the piece of information, starts the rumor. It goes around the world and comes back; in this case, it takes about six links, because it has to go from person to person to person. One of the ideas that came out of Watts and Strogatz and others is simply that judiciously placed links between people, random links even, that cross these boundaries will speed up the flow of information. Lots of people have showed results like this.

The more fundamental thing that came out of this whole line of work is that the exact pattern of the network matters: that by changing what the fundamental underlying network looks like, you can dramatically change the rate of information flow, who gets what information, and the kind of social outcomes you get. And this is just an example of a change speeding things up.

So first we said cognition matters. Then we found out that it's the exact nature of cognition that matters. We found from simulation that interaction matters; then we found that the exact pattern of interaction matters.

Chaos

Chaos has become a buzzword in organization theory [*slide 13*]. In the old paradigm, chaos is the natural state of the world, and the order is imposed. To get order, you have to have power, and you have to exercise that power. For example, the managers of an organization might dictate its culture. The new view of the world coming out of the simulation work is that order emerges naturally, and it emerges under a couple of very simple assumptions. One is that the agents have some cognizance of each other, that they have very simple models of each other. A lot of people — Kephart, Padgett, and many others — have work that speaks to this basic finding. The idea again is that knowledge of others enables realistic outcomes, and in particular it enables order. By "order" we're talking about everything from distribution of size of firm effects, to distribution of organization effects, to distribution of number of people in your networks, to overall patterns of coordination, and so on. To get these kind of ordering effects, your knowledge has to include a little bit of knowledge of others.

Chance and Paths

Another result coming out of this work is that chance, or the path you are taking, matters [*slide 14*]. The old paradigm included a variety of views. People would argue that your starting conditions should be irrelevant; that chance is irrelevant; that there is a single end state, and so on. Some of the work in simulation, however, suggested that this is just not the case. It suggests, in contrast, that most social systems are complex, that the systems are nonlinear, and hat therefore they exhibit path-dependency effects — so the order in which things happens, the starting conditions, and minor variations in starting conditions can have wildly different outcomes.

I like to think of this as a lost continent of research, because this is an old finding. It really dates back to some of the early systems work. Cyberneticists were saying this a long time ago. It was forgotten for a while, but now it's the big news again. If we put this result in the social context we are talking about, what it means is that it matters *who* you know *when*, it matters *what* you know *when*, it matters in what *order* you learn it. And those continue to matter over time. So it's not just having knowledge of others; the order in which you get that knowledge also matters.

Adaptation and Learning

The next finding is in the area of adaptation [*slide 15*]. Here a finding coming out of simulation is that if you're going to get adaptation, which we define as the ability to maintain or improve performance, you really need to avoid learning clashes. So far, we have been talking about agents that know stuff and are interacting, but we haven't really talked much about their ability to learn. There are many old paradigms that talk about organizational learning. For example, people talked about such things as, "Well, learning is the average of these experiences of the individuals in the company." That's experiential learning: how much you do something is important. In this case, turnover would be important, because since it's the average, it matters who's there when you calculate it. Or there is the argument coming out of population ecology that "organizations don't learn, they evolve." Or there is the argument coming out of formalization work, that learning is really imbedded in the storage procedures, databases, etc., within the company.

The newer view says, yes, all of those are right. It is not either/or; they are all right. In fact, we now have this ecology of learning mechanisms going on in the social world. There are many kinds of learning: experiential, structural, expectation-based, and so on. The psychologists say, "We've been telling you that for the past three decades." What that means from a social organizational process is that because there are many kinds of learning and because the learning is occurring not just at the individual level but at the synthetic agent or group level, those different kinds of learning can clash with each other. If they do clash, you won't get adaptation. So adaptations in part result from the absence of clashes.

You also need meta learning to respond to change processes, such as innovation, new technologies, and so on, and in many cases that means you need to have meta learning to learn how to trade off between exploration (that is, group or structural-level learning) and exploitation (that is, individual learning). A variety of models exist in this area, like the work by March and Leventhal on the exploration/exploitation model, our work on ORGAHEAD, and so on. But the basic finding from all of these is that learning clashes can impede adaptation.

Here is an example that shows how this effect came out in some simulations. We have a response surface showing a hypothetical organizational performance as a function of the organizational design. For simplicity, the only elements of the design I'm looking at here are size and density. The organization learns strategically and figures out what it wants to do and then decides what to do with it. This strategic learning moves the organization through this response surface by causing it to change its size and density by hiring and firing people, teaching them to talk to each other, putting them in change management seminars, whatever. The only thing is, the individuals themselves are all the while going through experiential learning, which has an S form for most human beings for most problems.

Now, let's imagine a particular organization that wants to improve its performance. How does it do that? Well, it decreases its size and tries to get people to talk together more, so it increases density, maybe by doing more team projects. Their actual performance, however, may not be this high. It may be much lower, because this is only the maximum possible performance. But why? Well, by making those changes — downsizing, getting rid of people, getting people to interact more, making them talk to each other — what are they doing? They are getting rid of the lessons of experience. They are obviating the value of old learning, and so their whole performance as a group is going to degrade.

We also have the performance over time of set of simulated organizations. Two things you will note. First, most of them are improving, on average. Second, the occupancy of first place changes sporadically. Now and then one of them accidentally really plummets. So if you explore why they plummeted, you will find in every one of those cases that it was because of a learning clash between one or more types of learning.

Transactive Memory

The next finding is in the area of transactive memory [*slide 17*]. This actually came out of computer science directly. Vegner and others said, "Look, in computer systems, a lot of the intelligence relies on the referential information and the ability of the computers to communicate with each other. It lies in the connections." He built a computer model of this and said, "Hey, maybe this is also true of human beings." Since then, there have been a number of experimental studies and field studies in transactive memory that have shown that in fact transactive memory is important, and it does improve performance in small groups, teams, organizations, and so on.

Transactive memory is your knowledge about who knows *who* and who knows *what*. The idea is that if you have models in which each of your agents not only has knowledge of others, but also has knowledge of who others know and what others know, they actually do better and the performance of the entire group improves. This is actually a very valuable thing for the social scientists to have found out, because it is a very important way of tying a historically important idea called social capital into our understanding of organizational performance.

Systems

The next finding is that that there are strong interactions between cognition and between interaction [*slide 18*]. In the old paradigm, you could talk about strict structuralism — that rocks and rats don't make a difference in how a network behaves, or there was psychological individualism — that is, individuals operate psychologically and don't care about what anyone else is doing. The new arguments are different. They say, "Look, not only is there an interaction between structure and people and cognition, there's actually more: what's really happening is that we have kind of soft boundaries between agency and structure."

To explain this idea, you have to understand the notion of the synthetic agent, which is an agent that's created out of other agents. So, for example, a group is a synthetic agent because it includes the individuals within it, as well as their knowledge of each other. Groups have the ability to learn, and they inherit all the properties of individuals in that sense. When you have synthetic agents along with human agents in a group, the boundary between what is an agent and what is a structure changes somewhat. As another example, consider that you as an agent have a certain set of knowledge at your disposal and a certain way in which you forget it and a certain way in which you use it. Let's imagine that you as an agent are busy working for Ford and you are painting those little racing stripes on the side of the cars. So the tools at your disposal include paint and the paintbrushes, and the car is completely distinct from you, the paintbrush is completely distinct from you, and there is a clear, physical boundary between you and everything else that's going on.

Now, however, Ford says, "Robots are going to be much better at painting these lines straight than human beings." So they replace the agent, the human being, with a new agent, which is the robot. But the brush and the paint can are built into the robot. So part of the old way — picking up the brush, dipping it in the paint — is gone. The boundary has been muted between agents. That's what we mean by mutable boundaries. The whole agent world is filled with discussions of mutable boundaries and how they can affect the outcome of group performance.

Some of the models in this area, like the e-commerce models, IBIZA, multi-agent models, or Soar docking with ELM, have shown that when you look at the interaction between a structure and agents, you see a huge interaction that totally influences the outcome. So you have to take both of them into account simultaneously.

One example comes out of some earlier work where we used Soar models and ELM (experiential learning model) agents and humans. We let all of them do a radar task where they had to determine for a bunch of planes going through space whether the planes were hostile or not. It turns out that if the agents are on a team or if they are in a hierarchy, the type of agent that performs best — the human, the Soar agent, the ELM agent — totally depends on both the task and the cognition. It is a complex nonlinear interaction between them.

SYNTHESIS AND NEW DIRECTIONS

So let's put all these things together. The findings from the past 40 years include that cognition really matters and that these limitations, the constraints on cognition and social action, affect the outcomes [*slide 19*]. Moreover, we found that the exact pattern of the information processing capabilities matters. It matters if your agents have interrupt schemes or not, for example. Interaction is critical to determining social outcomes, but it is not just interaction in general; it is the exact pattern in which those interactions occur. Knowledge about others enables more realistic outcomes and allows social order to emerge, and the order in which agents get that knowledge is also critical. The knowledge of others also has to include your knowledge of whom others interact with and what those others know in order to get effective performance or to get realistically human behavior out of these systems. Finally, when you move up to the group level, you get learning clashes that can impede adaptation and you get an interaction between cognition and structure. In a sense, that is a lot for 40 years, but what it really suggests is that we are currently seeing a new paradigm evolving in this area.

Now, let me illustrate what is happening because of this new paradigm. One thing is that people are now building new agent-based models that not only have increasingly smart and increasingly humanlike and increasingly cool agents, but these agents are put into networks. In fact, one of the big issues right now in robotics is how to build and think about and manage the networks, the communications patterns, between our robotic agents. If you go to any of the multi-agent conferences, they will talk about this.

Nonhuman Networks

The key thing here is that [social] networks are not exclusively among people [*slide 21*]. The network paradigm has permeated the area in terms of thinking about networks connecting agents to agents, but also networks connecting people to knowledge or agents to knowledge (the knowledge network), agents to tasks (the assignment network), agents to organizations (the work network), and so on. We have networks at every level connecting agents, knowledge, tasks, organizations, and so on. And that network-based approach is allowing modelers in this area to begin to say, "Oh, here's how my modeling hooks up to yours, because we are both representing our data in the same way." So there is an increasing amount of data sharing and model sharing.

Dynamic Networks

One difficulty, and the thing I haven't talked too much about so far, is that these little networks that we are talking about are dynamic [*slide 22*]. They start off in one way and then over time they keep changing, and eventually what you end up with is a totally new network. What you've got is a set of agents — little smiley-faced blue guys. They're in two different groups, with some interaction between them. Each of them has some knowledge. (I didn't show tasks.) And they are interactive. When they interact, they learn from each other, that is they communicate and send new information, which leads to the growth of new information at those purple explosions, and new people come into the company. And of course the groups reformulate and change. So, for example, this guy who was connected over here learns new things that makes him more like this group, so eventually he might move over to this side, taking his knowledge with him.

One of the big issues in this area is group emergence; that is, how do I know when a group emerged, how can I track it, how can I understand it, how can I see it? That's an unsolved problem. Simon has been telling me for years that we need to solve this. Yes, it's true. It's not solved.

Technology and Networks

Another problem in this area is dealing with technology. I told you we can treat technology as an agent. That's fine. However, the big issue is not just treating technology as an agent; it's the fact that when we talk about technology, we are talking about new information technology giving new access to more people, more information, more of the time. So the issue is how to represent this situation in such a way that we can look at matters of scale. This question is particularly important if you're trying to model, say, the whole Internet, for example. So people are starting to use other techniques for looking at really large-scale agent networks.

Here's another example of a result from this area. This is a transactive memory model that looks at organizational performance over time when you've got worlds that have just people or people plus databases or people plus avatars in the company. One of the things you see happening is that when you include databases in these systems, they slowly decrease the amount of know-who, whereas if you have avatars, people know more about who's out there, and avatars also decrease what people know. More importantly, the other thing that technology does is change the fundamental way in which information is shared, the extent of that sharing, and the extent to which people interact. So, for example, if you put a Lotus Notes database system in your company, one of the big impacts will be that it will actually reduce the overall level of shared knowledge in the company, delay people interacting with each other, and lead to less interaction overall. These changes may not be that great from the perspective of living in the company, but they may have no impact on corporate performance.

Applications for Enhanced Networks

I want to close by extending that point in the organization setting and telling you about a real-world application. For this application, we went out and collected network data from a real company — it happened to be the Navy. We plugged all this into the system and used it to predict things about organizational adaptation.

One of the first predictions was that to be a really successful team, what they should do is spend more time retuning or retasking people and engaging in redesign tasks — that is, changing who is reporting to whom and who is doing what. If they spend time hiring and firing their personnel, they're probably going to be low performers. That was a generic finding coming out of this work. We then went and collected data on an actual platoon who were playing this one particular war game, plugged it into the model, and predicted who would be their emergent leaders. We got that exactly right, because we predicted that their emergent leaders would be those who had a high workload and who preferred to shed tasks rather than accept new personnel, which, of course, they did.

But we noticed something very interesting in that experience. In their after-action report the participants would talk to each other about why they did what they did. We just asked them, "Well, do you know what each other is doing?" Let's say I was the guy who was trying to shed a task. I'm Bravo. I've got a high workload, I'm trying to give you my task. You're not necessarily accepting it. So you're trying to do it, but you're not accepting it. So afterwards, we go to the other guys, Alpha and Charlie, and say, "Well, what do you think Bravo was trying to do?" They say, "I don't know. He was just sloughing off. I have no idea why he was asking me to go and do this task." No idea. They have no common view of the world.

Then we went to the admiral and said, "We understand that you're running the simulation for us, but could you have these guys tell each other what they're doing?" So the next time, they began by telling each other what they were doing. The next time the guy who was the emergent leader was trying to shed tasks, other people accepted them, because they knew what he was trying to do. So we thought, "This is really interesting." It fit our model, but we also saw something that we didn't know would happen, namely, that there seemed to be a trade-off between having a common view of the world — a common operational picture, a common understanding of what others were doing — and the performance of the group. We designed these groups to have no communication and to have optimal performance, and they worked. They were high-performance teams, they had no communication, they were great. They also had no view of what the others were doing.

So then we went back to the simulation and said, "Will this observation that we just had come out of the model?" The results of simulations that were done after the fact showed that the factors that led to high performance were dramatically different than those that led to having a common operational picture or having high adaptivity. My point about this is that the new work in this area is integrating data much more heavily. There's a real synergy in the new models because of the network perspective in going back and forth between real data and simulations.

This new approach is being used in a whole bunch of areas; these are just some of them [*slide 29*]. Each of you has many, many others. But the point I want to make here is that in this area there are applications both at the research level and at the totally applied level. I mean, how much more applied can you get than knowledge management, human resources policy evaluation, and looking at bioterrorism — bioterrorism isn't up there; that's a new project. And of course there are many universities, as you know, where this is being done. This is just some of them with big programs in the area, and for the graduate students here, I'll be glad to share this list with you.

Computational Social Science: Agents, Interaction, and Dynamics

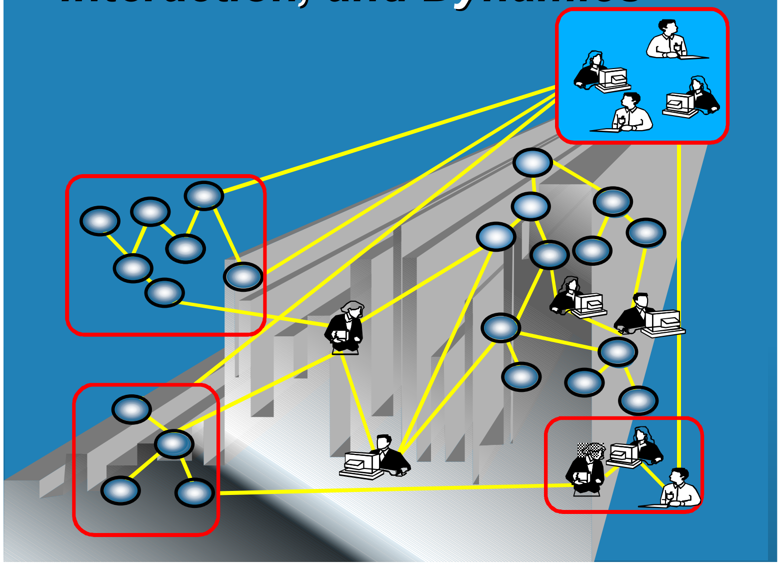

Carnegie Mellon

CASOS

Center for Computational Analysis of Social and Organizational Systems

Kathleen M. Carley

Director & Professor
Carnegie Mellon University
Porter Hall 219a
Pittsburgh, PA 15213
Tel: 1-412-268-3225
Fax: 1-412-268-6938
Email: kathleen.carley@cmu.edu

2000

Computational Techniques

- Important tool for social theorists and methodologists.
- Multi-agent models particularly valuable.
- Enable the researcher to:
  - Examine the relations among groups and individual agents.
  - Witness the way in which these relations enable, constrain, and affect agent and group behavior.
  - Examine emergent phenomena.
- Key areas:
  - Nature of coordination.
  - Facilitating or inhibiting change.
  - Effect of scale.
  - Effect of technology.

2

The following slides have been omitted: slides 10, 16, 20, and 23-28.

40 Years of Computational Theory

- Computational analysis aided in revolutionizing social theory.
- Important advances in the areas of:
  - Cognition.
  - Interaction.
  - Chance.
  - Adaptation.
- Fundamental results:
  - Bounded rationality.
  - Satisficing.
  - Competency traps.
  - Emergent order.
  - Learning clashes.
- New paradigms such as information processing theory.

Question existing paradigms
Explore non-standard assumptions

3

Illustrative Models: Black-High, Star-Medium, and White-Low

	Agent		Task			Structure		Resource	
	Learning	Multi-Agent	Types	Reassignment	Rework	Types	Dynamic	Kind	Amount
VDT	❍	◗	❍	◗	●	❍	❍	◗	●
TAEMS	❍	●	●	❍	❍	◗	❍	●	●
STEAM	◗	●	◗	❍	❍	❍	◗	◗	●
COMIT	●	◗	◗	◗	❍	❍	❍	●	●
TacAir-Soar	◗	❍	❍	❍	❍	❍	❍	❍	❍
Orga-head	●	◗	❍	●	❍	●	●	◗	●
Garbage Can	❍	◗	❍	❍	❍	◗	❍	❍	●
NK	❍	◗	❍	❍	❍	❍	❍	❍	◗
Construct	◗	◗	❍	❍	❍	◗	●	◗	●
ELM	◗	◗	❍	●	❍	●	●	◗	●

4

The Many Faces of Models

- Models as theory
  - The model does the task it seeks to explain
  - Social Turing test
- Models as agents
  - Tournament participants
  - Substituting for humans in the lab
- Models as virtual world
  - Social theory from the ground up
  - A-life
- Models as empirically grounded theory
  - Empirical data as integral part of model
  - May be ethnographic
- Models as hypothesis generators
  - Response surface analysis
  - Systematic generation of hypotheses

5

Computational Analysis Approach

- Multi-agent models
  - Agents are humans or artificial
  - Collections of diverse agents
  - Heterogeneity of agents
  - Social action derived through interaction
- Socio-information processing approach
  - Social linkages, constraints
  - Cognitive capabilities, constraints
  - Technological changes, constraints
- Emergent behavior
  - Individual, group, organization, population level
  - Patterns
  - May not be equilibrium
- Empirically grounded
  - Models based on empirical findings
  - Parameters, mechanisms, input may be from other data sets
  - Multi-level validation

6

The Nature of the Social Agent

- Fundamental question
- Computational analysis generated many of the puzzle pieces
- Interdisciplinary
- Teams of scientists

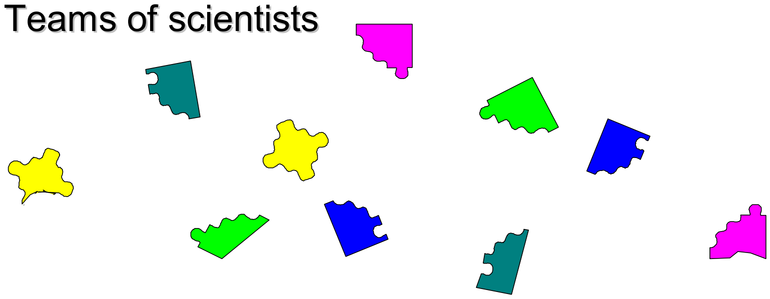

7

Cognition: Constraints Matter

- One of earliest social science computational models
- Old paradigm - *rational actors*
- New paradigm - ***Boundedly rational actors***
- Models:
  - A behavioral theory of the firm - Cyert and March
  - Garbage can model - Cohen, March and Olsen

Cognitive and social limitations affect outcomes

8

Cognition: Architecture Matters

- Alternative paradigms - agents are agents
  - Strong structural position
  - Rational actor - agent differences not relevant
  - Decision making - outcomes based on decision process, not information process
- New paradigm - **agent variation**
  - Agents are heterogeneous
  - Technology as agent
  - Information processing capabilities matter

Exact information processing capabilities affect outcomes

9

Reciprocity - Interaction Matters

- Game - Prisoner's Dilemma
- Issue - what type of social intelligence is needed to get groups to cooperate?
- Tournament - Axelrod
- Outcome
  - Tit-for-tat
  - Reciprocity norm

0 intelligence
Order effects
Alternate norms
Alternate choices

Interaction is a critical determinant of social outcomes

11

Networks Matter

Watts and Strogatz
Young
Markosky
Many others

Small world phenomena

The exact pattern of the network matters

12

Chaos: Knowledge of Others Matters

- View - Leviathan - chaos as the natural state
- Old paradigm - order is imposed
  - Requires power
  - Requires the exercise of power
  - E.g., organizations - managers dictate culture
- New paradigm - **order emerges**
  - Agents need to be cognizant of others
  - Agents need to have models of others

Kephart
Padgett
Many others

Knowledge of others enables realistic outcomes

13

Chance: Path Matters

- Old paradigm
  - Starting conditions are irrelevant
  - Chance is irrelevant
  - Single end state
- New paradigm
  - Complex systems
  - Social systems as non-linear
  - Path dependence
  - Minor variations in starting conditions can have wildly different outcomes

The Lost Continent
Forrester
Sterman
Many many many others

Who you know, what you know, what you learn when is critical

Adaptation: Maintain or Improve Performance

- Old paradigms - organization learning?
  - Average of agents: turnover, experiential
  - Organizations don't learn - they evolve
  - Formalization - procedures, databases, etc.
- New paradigm - learning ecologies
  - Many kinds of learning
  - Change processes - innovation, technology, legislation
  - Adaptation
    - Absence of clashes
    - Meta learning
    - Trading off exploitation and exploration
- Models
  - EE - March, Levinthal
  - ORGAHEAD - Carley

Learning clashes can impede adaptation

Transactive Memory

- Idea:
  - In computer system, intelligence lies in part in connections and referential data
  - Maybe this is true of human cognition
- Transactive memory
  - Your knowledge of
    - Who knows who
    - Who knows what
- Models
  - TM model - Wegner
  - Construct-o - Carley

Who knows who knows what and who knows who affects outcome, TM improves performance

17

Interactions: Systems Matter

- Old paradigm
  - Strict structuralism
  - Psychological individualism
- New paradigm
  - Soft boundaries between agency and structure
  - Synthetic agents
- Models
  - E-commerce, IBIZA - Krishnan
  - Multi-agent soar docking with ELM – Carley, Lin, Prietula

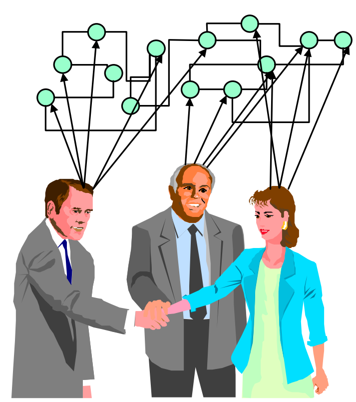

Interaction between cognition and structure influences social outcomes

18

Putting the Pieces Together

- Cognitive and social limitations affect outcomes
- Exact information processing capabilities affect outcomes
- Interaction is a critical determinant of social outcomes
- The exact pattern of the network matters
- Knowledge of others enables realistic outcomes
- Who you know, what you know, what you learn when is critical
- Who knows who knows what and who knows who affects outcome, TM improves performance
- Learning clashes can impede adaptation
    - Interaction between cognition and structure influences social outcomes

New Paradigm

19

A Meta-network Representation of Organizations

	Agents	Knowledge	Tasks	Organizations
Agents	**Interaction Network** *Who knows who*	**Knowledge Network** *Who knows what*	**Assignment Network** *Who does which tasks*	**Work Network** *Who works where*
Knowledge		**Information Network** *What informs what*	**Needs Network** *What is needed to do which task*	**Competency Network** *What knowledge is where*
Tasks			**Precedence Network** *Which task precedes which*	**Market Network** *What tasks are done where*
Organizations				**Inter-Organizational Network** *Which organizations interact*

21

Dynamic Networks

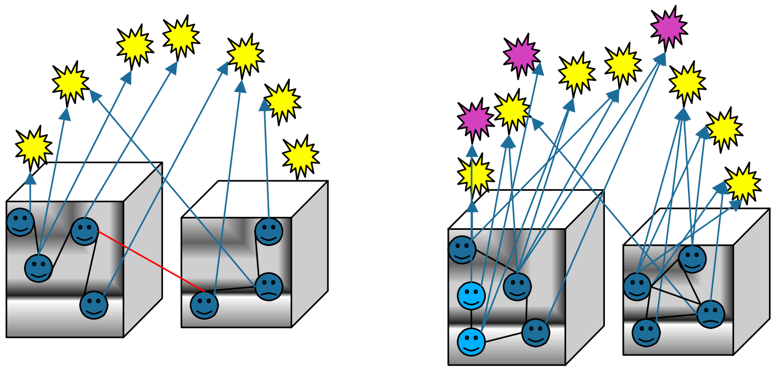

22

Application Areas

- Team design
- Organizational architecture evaluation
- Team flexibility and robustness analysis
- Tech transfer schemes
- HR policy evaluation
- Knowledge management
- Interaction of training and design
- Vulnerability analysis and information security
- Impact of IT
- Social network evolution and illicit drug/information transmission
- Risk management vs. information diffusion vs. early disease/virus detection

29

Academic Interest

- Universities and educational/research units forming programs
  - Stanford
  - Michigan
  - MIT
  - University of Texas - Austin
  - University of Texas - Dallas
  - John Hopkins
  - Cornell
  - Duke
  - University of Rochester
  - University of Arizona
  - UCLA
  - Penn State
  - University of Chicago
  - University of Illinois Urbana Champaign
  - Santa Fe Institute
  - England
    - University of Manchester
    - University of Surrey
    - London School of Economics
  - The Netherlands
    - Groningen
    - University of Amsterdam
  - Italy
    - University Trento
    - University Degli
    - University of Bologna
  - France
    - Insead
  - Japan
    - Kyoto University
    - University of Tsukuba

DISCUSSION: OPENING PRESENTATION*

T. WOLSKO, Argonne National Laboratory, Moderator

Michael North: I'm Michael North from Argonne National Laboratory. One of the things we've been talking about here that I think is very interesting is the idea of increasing specificity, in the sense that people are going from a general notion of finding an optimal solution to realizing that the optimal, if it exists at all, depends very much on the exact situation. I think the Lotus Notes example is particularly relevant to what we're talking about. Clearly, in a lot of cases, as you're seeing, it may make things worse. But if you have certain types of organizations — ones with high turnover, for instance — Lotus Notes may actually help. Could you elaborate on some of the things that you've seen change in people's thinking about moving from a general optimum to very contact-specific optimum?

Kathleen Carley: One of the main things I've seen happening is that a lot of the models today are much larger, much more detailed and specific than earlier models. That makes it difficult to publish these models in the journals; there are just too many pages of codes. So we have some issues in the field about how to archive these models — how can we share them? And that's led to a whole new research area on what is the right infrastructure for model sharing.

Because these details really matter and because there are so many findings out now showing that as you change, say, the cognitive architecture or the interaction, you change the results, you see people paying more attention to trying to dock or compare their models one against another. We don't have something like meta-analysis in psychology, yet there is more of this kind of discussion. In some sense, it's like this area is out of its infancy but is still in its adolescence.

North: Yes, that's one of the things I was going to note, too, in terms of repeatability. That's a hallmark of the scientific method. But how do we try to set up systems that are repeatable, given that not only are the problems so specific — you can slightly vary things and get a new problem space — but also that all the tools we're using are so chaotic and so complex in how they depend on those initial conditions.

Carley: Well, one of the answers that a lot of universities are looking at is open-source code. And there's another approach, which says that it doesn't really matter if your models are identical. We do a type of meta-analysis, and we get the same basic finding from a set of models, and that finding is very robust.

So that would be another possibility. The third thing, though, is some new work going on by logicians who are actually trying, on an experimental level, to show that, given an experiment that's completed — a virtual experiment on the computer — can you then back-prove that it occurred logically from its underlying assumptions? That's very state-of-the-art.

* [Editor's note: The discussion sessions were recorded with the speakers' knowledge and transcribed. The transcripts were edited for continuity and ease of reading; every effort was made to identify speakers and interpret comments accurately.]

North: Then this experimentation just seeds ideas to a certain extent.

Claudio Cioffi-Revilla: Claudio Cioffi from the University of Colorado. Could you elaborate on the notion you mentioned that "how groups emerge is an unsolved problem." Surely you meant there may be new *solution*s to how groups emerge? There's a game theory of how groups emerge; there's a stochastic theory. What did you mean?

Carley: I'm glad you asked that. What I'm talking about is not the process by which groups emerge. We have lots of theories, lots of work on that topic. The problem is, in an agent-based world, I want to be able to see a set of agents and say, "Aha, that's a conglomerate; there is a group." Some work has been done on recognizing when a set of agents is a group using clustering techniques, divisionary techniques, and so on. But if a mechanism is not predefined, what there isn't right now are good automated techniques for saying, "That agent belongs to this group and not that group."

It's not the *process* of groups evolving, it's recognizing that what I have is a group. A lot of the best work right now still involves drawing up a diagram and saying, "Yes, I can circle it; that's a group." Or you recognize a group, then backtrack, analyze it statistically, and say, "Okay, it kind of fits with my cluster analysis, it kind of feels right, so it's a group." So it's coming from that level, not from theory.

Jonathan Bendor: Jon Bendor of Stanford University. Going back to your remarks about the importance of starting positions and the anything-can-happen point of view. I'd like to generate a discussion about that, because I come from a somewhat different research tradition. In the research tradition that I belong to, assuming anything can happen is generally a weakness of a research program, because it reduces the empirical content of theories. If anything can happen, you look at a particular phenomenon and ask, "Why did this happen?" and you answer, "Well, give me the right starting position and I'll give you that outcome." Then the theory might well be vacuous — it doesn't exclude anything from happening. And a theory that doesn't exclude anything from happening is unfalsifiable.

It seems to me that in quite a few of these theories, however, if you work at them a little, you can show that there will generally be a unique limiting distribution to them. That doesn't mean that you reach a static kind of equilibrium — the individual agents or the individual populations of agents don't eventually stop changing. What happens, though, is that you reach a probabilistic equilibrium, in the sense that a Markov chain, for example, generally reaches a unique limiting distribution — and that's a lot of the empirical content of the theory. So you're saying, "Probabilistically, we know that it's going to go to a steady state, and we're going to try to characterize some of the properties of that steady state." Along the way, of course, the path matters. It's going to look different if you start here versus somewhere else, but in the long run you are going to reach a probabilistic steady state.

Carley: You're overinterpreting what I said. The issue is not that given your different starting positions you can end up anywhere, it's that given different starting positions you may end up in different ending conditions. Most of these models don't have the feature that you can end up anywhere. They do have sets of outcomes at some point. But you will get two very different ending conditions — not all possible endings, but different ones, depending on where you start.

Bendor: Well, what is meant by an ending condition?

Carley: Take one of our organizational adaptation models. If you run this model for several time periods, you see a behavior pattern where you get a probabilistic average of always hitting the same areas — just like what you were talking about. But there will be two different probabilities. If you look at the data, they will bifurcate. One set of organizations will end up with a high average performance level, another set will end up with a low average performance level — and where they started will affect where they end up.

North: In the work we're doing, end states and these long-run issues do matter to some extent, but we're actually much more concerned with the *process*, the path. The fact that these distributions may exist doesn't make a whole lot of difference to us. We care about the path of the market through the space. We also feel that in many cases the long-run equilibrium, or statistical equilibrium, will never really occur, because by the time you get that far out the whole environment has changed. So for us, the existence of these distributions is nice, but they may not actually matter.

Bendor: Two comments on that. First, there are results on the rate of convergence, and often it is surprisingly fast. And back to your first point, if you know what the long-run steady-state distribution of the system is, that gives you some information about the path.

North: The problem is that just knowing the starting and ending points is not enough. The path could do anything in between. A great example is the military.

Bendor: Anything with high probability?

North: We're trying to say that low-probability events can make a big difference here. A good example is a large military force fighting a small military force — there have been studies that indicate the large military force can lose. It's low probability, but it matters to the loser.

John Bower: John Bower from London Business School. I wonder if you'd comment a bit more on the techniques that people are starting to use to analyze the output from agent models. I get the strong feeling, in fact a slightly queasy feeling, that we often try to convince people about the output and the quality of the models by saying, "It quacks like a duck, so it's definitely a duck."

Carley: If you're in artificial intelligence, that's exactly the test that's used, because you only need to show proof of concept. But let me back up a second. A whole variety of new techniques are being developed for analyzing output from agent models. Some of these are standard statistical techniques, but there are also new techniques for network-based models. A new set of techniques derived from time-series analysis lets you look at complex data over time. There's also work coming out of engineering; many people use Fourier transforms, for example. But it is still the case for much of the work out there — and this is something we have a lot of trouble with in the journals — that people will run a model once and present their results as definitive, when in fact the model has stochastic elements that have to be demonstrated over hundreds of runs.

The other factor is that we have very large models, with huge numbers of agents in them. They actually break most standard statistical packages. If you're looking at networks of agents, the models will break all the network statistical packages. And if you want to visualize the results, they'll break all the visualization packages. One of the big research areas being funded by NSF right now is in designing analytical software that can handle these large models.

Robert Reynolds: Bob Reynolds, Wayne State University. I noticed that in your summary of results you didn't specifically address agent organization and problem-solving, though you talked about the infrastructure in which agents interact. For a particular problem, or an environment that poses a sequence of problems, how do the agents adapt?

Carley: There's some work in this area, but not a lot that has led to coherent findings. Some findings have shown that agent structures can adapt to a task environment. There's work that shows that when agents are structured in a way that's perfect for the task, they will perform better. We know how to design teams for tasks now; that's a pretty resolved area at this point. And we know about methods of adaptation. But in terms of what types of adaptations are needed for what types of tasks — other than some findings about teams, hierarchies, and volatility of tasks — we're still learning.

Artificial Markets

FINANCIAL MARKET EFFICIENCY IN A COEVOLUTIONARY ENVIRONMENT

B. LeBARON, Brandeis University*

ABSTRACT

What does the evolutionary interaction between different types of investors look like? This paper discusses research trying to understand the evolutionary dynamics between agents using differing lengths of past data to make decisions on portfolio choice. Computer simulations of a simple agent-based model show that agents taking a long perspective on past data have a difficult time dominating shorter perspective agents. The resulting dynamics replicate many features of actual markets. Furthermore, strategies become more homogeneous near sharp price declines, suggesting a liquidity-based explanation for market crashes and excess volatility.

INTRODUCTION

Traditional equilibrium models for financial markets often rely on the process of evolution, either explicitly or implicitly. These models generally assume that asset prices are the result of market participants holding common rational beliefs about the behavior of economic variables in the future and acting in a well-prescribed fashion on these beliefs. These economic worlds are simple, tractable, and generally at odds with the empirical evidence. Observed financial markets often appear too volatile and too predictable to be explained as the outcome of well-learned rational beliefs and strategies. Furthermore, the existence of large trading volume in financial markets adds important questions concerning heterogeneity across participants, since if all people agreed on asset valuations, trade would be unnecessary in most situations. This paper uses an agent-based model of a simple financial market to explore the evolutionary aspects of market dynamics, with the goal of understanding the barriers to market efficiency that cannot be eliminated through evolution and learning alone.

Most modern analysis of financial markets includes two crucial assumptions: markets are populated with rational agents, and they are in some kind of stationary equilibrium. Together these two assumptions yield tractable, testable restrictions for well-crafted theories. The first assumption can be relaxed somewhat in that irrational players may appear at times, but their suboptimal strategies will be driven out of the market.[1] Weakening the second assumption causes

---

* *Corresponding author address:* Blake LeBaron, Graduate School of International Economics and Finance, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453-2728; e-mail: blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a faculty research fellow at the National Bureau of Economic Research and an external faculty member at the Santa Fe Institute.

[1] Friedman (1953) provides the most often cited arguments for an evolutionary foundation for assuming rational behavior. In his case, for flexible exchange rates, he specifically comments on how less-than-rational speculators will be driven out of the market. Recent research on noise trading (DeLong, Shleifer, Summers & Waldmann 1991) and evolution (Blume & Easley 1990) has begun to suggest some flaws in the evolutionary argument for rationality in a financial setting.

more difficulty, since it is closely linked to the rationality assumption. Out of equilibrium, it becomes difficult to judge rational versus irrational strategies, as the economic landscape is in a continuing state of change. The biological term for this is coevolution. Strategies evolve against this current set of strategies in the population and not some well-defined fixed fitness norm.[2] In such a world, rationality can only be judged relative to the current population and not some well-defined fixed target that players should want to attain. It would be convenient to argue that these out-of-equilibrium dynamics can be ignored. However, this leaves open the critical question of how markets reach equilibrium in the first place.

In order to analyze out-of-equilibrium dynamics as well as convergence properties, markets will be populated with boundedly rational learning agents.[3] These are relatively simple agents, trading and learning about price dynamics as they go along. This facilitates the analysis of overall market dynamics and convergence properties when agents are faced with the same situations seen by ordinary people. It is important to realize that boundedly rational does not necessarily translate into stupid agents. They are often faced with situations where being completely rational may be computationally intractable, involving the beliefs of all the other market participants along with their dynamic decision-making processes. The only option available is to follow simple rules of thumb that are empirically tested and adjusted over time through learning.

This paper focuses specifically on one type of rationality, the appropriate use of past information. If financial time series were completely stationary, then more historical data would always yield better investment decisions. However, in practice, it appears that many market participants chose to ignore some past information to focus on the present. Recent arguments about a "new economy" are a good example of this. This behavior might indeed be rational if markets have changed, yielding the historical data irrelevant and giving those following it the survival chances of dinosaurs.

An evolutionary struggle between short- and long-horizon investors will be explored in an attempt to assess when and if the long-horizon types will evolutionarily dominate the market. The setting will be one with a completely stationary dividend process. In such a world, it might seem obvious that the long-horizon investors should dominate, but this is not necessarily the case. Prices move endogenously according to the traders' strategies and can even move in such a way as to enhance the strategies of short-term traders. It is also important to realize that in such a market, it is not clear what is a rational or irrational strategy without the guidance of a market safely in equilibrium. Given a turbulent market of short-run investors, it may be individually rational to become one of them, as opposed to taking the more difficult path of sticking to a long-run perspective.

---

[2] A good biological example is to think about evolution against predators versus evolution against climate. In the later case, one is probably safe in assuming a fixed fitness landscape, but in the former, this landscape is an ever-changing target.

[3] The concept of boundedly rational agents was introduced by Simon (1969). Recent applications in macroeconomics are summarized in Sargent (1993) and Sargent (1999). Often the argument for bounded rationality rests on actual bounds on computing power in the brain. However, bounded rationality might be argued for in terms of robustness. In a complex financial world, all strategies will be incomplete in some aspect, so it can be argued that simpler strategies may do better in terms of avoiding some really big mistakes. This is a little like arguing that you don't want to get "too smart for your own good."

MARKET DESCRIPTION

The market simulations used here are part of the class of economic models referred to as "agent based." Models of this type consist of large numbers of interacting agents, each acting independently of the others, often with active learning and adaptation.[4] Agent-based markets share many features: many interacting individuals, evolutionary dynamics, learning, and bounded rationality. However, the key distinguishing feature is that heterogeneity itself is endogenous. Markets can move through periods that support a diverse population of beliefs and other periods where these beliefs and strategies might collapse down to a very small set.

The market is a very simple one with a single equitylike security paying a random dividend each period and available in a fixed supply of one. This dividend follows a stochastic growth process that is calibrated to aggregate dividend series for the United States. There is a risk-free asset that is available in infinite supply paying a constant real interest rate of one percent per year. Portfolios are rebalanced and trades are made at a monthly frequency. Also, prices are determined and dividends are paid each month, which can be thought of as the basic unit of time in the market. Therefore, this is more of an experiment concerned with longer-term macroeconomic behavior as opposed to the minute by minute dynamics of day trading.

The basic actors are a set of 1,000 agents. These agents adjust their portfolios and trade independently. They have well-defined objectives in terms of optimal portfolio allocations.[5] However, they differ in one key respect: they have different views about how much past data are relevant in making their decisions. Some may take a long-horizon perspective using an equivalent of the past 20 years of data, while others view that only the past year or two of data are important. For them, previous data have become irrelevant in the investment decision-making process. As trades and time go by, the agents accumulate wealth and consume. Evolution takes place by eliminating agents with the lowest amount of wealth and replacing them with new ones. In this way, a "survival of the fittest" dynamic is imposed on the population of trading agents. As mentioned previously, it is important to understand how much pressure this puts on the market to move to a homogeneous rational outcome.

An important piece of the market is given by the trading strategies themselves. These can either be thought of as rules of thumb followed by the investors or as institutional managers that dynamically adjust their clients' portfolios. Since there are only two assets in the market, these strategies can only be market timing strategies if they deviate at all from simple buy and hold strategies. The strategies convert a subset of current market information into a recommended

---

[4] Examples of this include Cont & Bouchaud (2000), Epstein & Axtell (1997), Palmer, Arthur, Holland, LeBaron & Tayler (1994), Arthur, Holland, LeBaron, Palmer & Tayler (1997), Kim & Markowitz (1989), Levy, Levy & Solomon (1994), Lux (1997), Tay & Linn (2001). Also, the website maintained by Leigh Tesfatsion at www.econ.iastate.edu/tesfatsi/ace.htm is an important source for agent-based research in economics. Finally, the site at www.brandeis.edu/~blebaron/acf summarizes agent-based research in finance, and a survey of some of the early research can be found in LeBaron (2000). Some commentary on the construction of agent-based models is given in LeBaron (2001a).

[5] The agents have logarithmic preferences over expected future consumption. Their time rate of discount is set to 0.95 per year. This is a well-understood optimization problem in economics and finance (Merton 1969 and Samuelson 1969). It yields a consumption value that is a constant fraction of wealth and an investment strategy that should maximize expected log returns.

portfolio allocation. The allocation gives a fraction of savings to put in the risky asset.[6] The market information includes past returns, dividend yields, and two moving average technical indicators. The rules can be built off this information in any combination.[7] It is also important to realize that rules can ignore any piece of information too. In many ways, the ability to ignore superfluous information is part of what is being tested.

Two further issues related to trading strategies remain. First, agents must decide which rules to use. At any given time, there is a set of 250 active rules (investment advisors). Agents choose the rule that has performed the best in terms of their objective function. In doing so, they make their decision based on past data using only what they feel is relevant. In other words, if the agent believes that the only the past two years of data are important (a relatively short horizon type) this is the range over which available strategies will be evaluated.[8]

The second crucial issue is how rules learn and adapt over time. If an investment advisor has at least one agent signed up for its services, it will continue to exist with no change in its interpretation of market information into a dynamic strategy. If the advisor finds itself with no customers, it will be eliminated and replaced with a new advisor. The new advisor is created from the current population of active advisors using a genetic algorithm. This gives an interesting evolutionary dynamic to trading strategies. Those that are being used survive, and those that aren't are eliminated. The genetic algorithm tries to bring useful features of the current active strategies in to future ones. Success is determined purely on whether anyone is using a given strategy.[9]

Trading takes place each period. Agents all enter the market equipped with a chosen rule and their current portfolio positions. This gives a well-defined function for shares as a function of any given market price. Therefore, in principle, the market could be cleared by a Walrasian auctioneer operating each period. This is essentially what is done. A numerical procedure is used to find a price that sets the demand for shares of the risky asset equal to the fixed supply in the market of one share.[10]

COMPUTATIONAL EXPERIMENTS

Benchmark Runs

[6] Before making asset allocation decisions, agents consume a certain fraction of wealth, leaving the rest for investment. Also, they are not allowed to borrow or to sell shares short.

[7] The actual structure of the rules is given by an artificial neural network and is detailed in LeBaron (forthcoming 2001b). These provide a flexible functional form for turning the data inputs into actual portfolio weights.

[8] To introduce further heterogeneity in trader behavior, it is assumed that traders only evaluate a small number of rules each period. They do not do a complete search over all possible advisors. Specifically, their current advisor is compared to one chosen randomly from the upper half of the advisor distribution measured using the agent's own view of how much history to use.

[9] It would be difficult to use any other fitness measure, such as expected returns, since agents don't view these from a common perspective. They are estimating returns over differing time horizons.

[10] A more realistic market might consider the actual market microstructure. However, since this market will be viewed as a fairly long-range (monthly) pricing series, the temporary equilibrium assumption does not seem unreasonable.

The following sections provide example runs of the market. In all cases, the dividend series follows a random walk that is roughly calibrated to postwar U.S. aggregate dividends with an annual growth rate of 2 percent and an annual standard deviation of 6 percent.[11] The risk-free rate of interest is fixed to a constant 1 percent per year. All these rates are adjusted to monthly frequency, which is the benchmark time horizon for trading and dividend/interest payments.

The key variable of interest in these experiments is the horizon length of the agents. This represents the distance they look into the past. Two different experiments will be considered. The first, referred to as *all horizon*, uses agents drawn randomly from 5 to 250 months in length. This allows for a diverse population with many different investment horizons competing against each other. The second experiment, referred to as *long horizon*, loads the market with a set of agents using a relatively long time horizon. In this case, agents are drawn from a distribution between 220 and 250 months. This loads the market with only long-horizon investors. The objective is to see if this group performs differently from the diverse investor horizon case.

The market is run for 10,000 periods, which, in calibrated time, is actually over 800 years. Figure 1 shows a run for the first 3,000 periods for a set of all horizon agents. The figure displays both price and volume. The figure shows three different phases for the market. In the first, the market is slowly adjusting, and the price is catching up to get on the correct exponential growth path. In the second phase, the market appears almost cyclical in its fluctuations about the constant growth trend. There appear to be smooth cycles that don't look very reminiscant of actual markets. In the final phase, after period 1500, the market begins to look more normal, with some run-ups in the price followed by some sudden crashes. Trading volume also increases during this later period too.

Figure 2 presents the same information for the long-horizon investors. Here we see a very different picture. After an initial adjustment phase, the price maintains a relatively steady growth rate, and trading volume drops to near zero.[12] This picture shows a glimpse of a market more closely resembling a traditional efficient market equilibrium. Further results will reinforce these early pictures.

In Figure 3, a detailed picture of the two price series is compared. This figure displays prices taken from the final 1,000 periods (roughly 80 years) of a 10,000-period run. This is done to capture behavior after all agents have settled down in their learning periods. The first two figures demonstrate that during the early periods, behavior might not reflect the final outcomes in the different market experiments. This picture very clearly repeats the message of the early figures. The top panel, which corresponds to the case where traders are using diverse horizon lengths, shows a market price that occasionally takes some large swings, often ending in dramatic crashes. In the lower panel, where traders are only long horizon, the picture shows a much smoother price dynamic with fewer large moves. The price still does indeed move, which reflects the fact that the fundamental dividend series is a somewhat volatile random walk. However, the amplification of volatility from the lower to the upper panel is clear.

---

[11] See Campbell (1969) for a summary of values for several different countries.

[12] Investors enter the market with no innate knowledge of how it works, or how prices move with the fundamental, so it is sensible that some learning must take place for a short time.

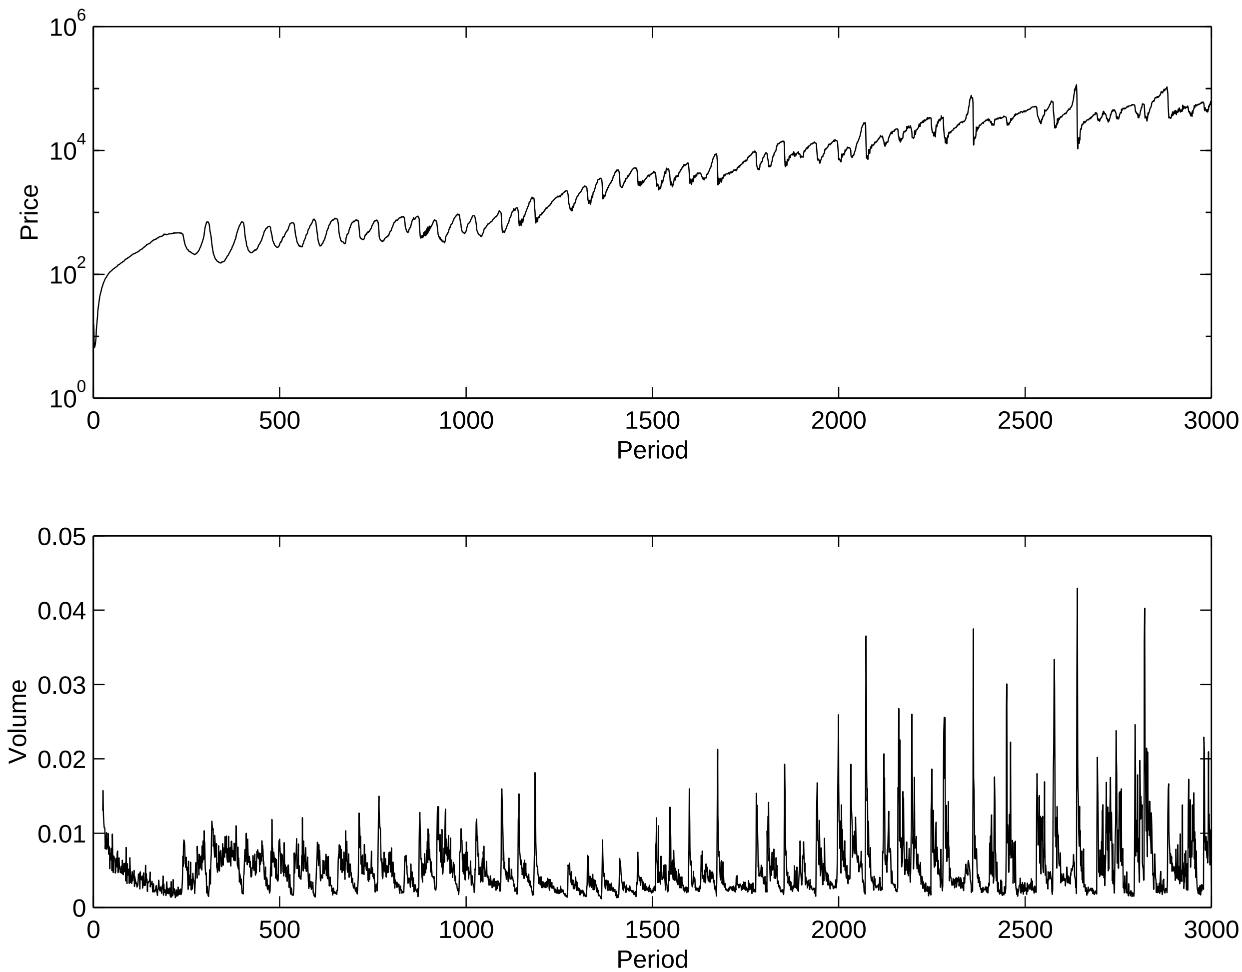

FIGURE 1 Price/Volume: All Horizons

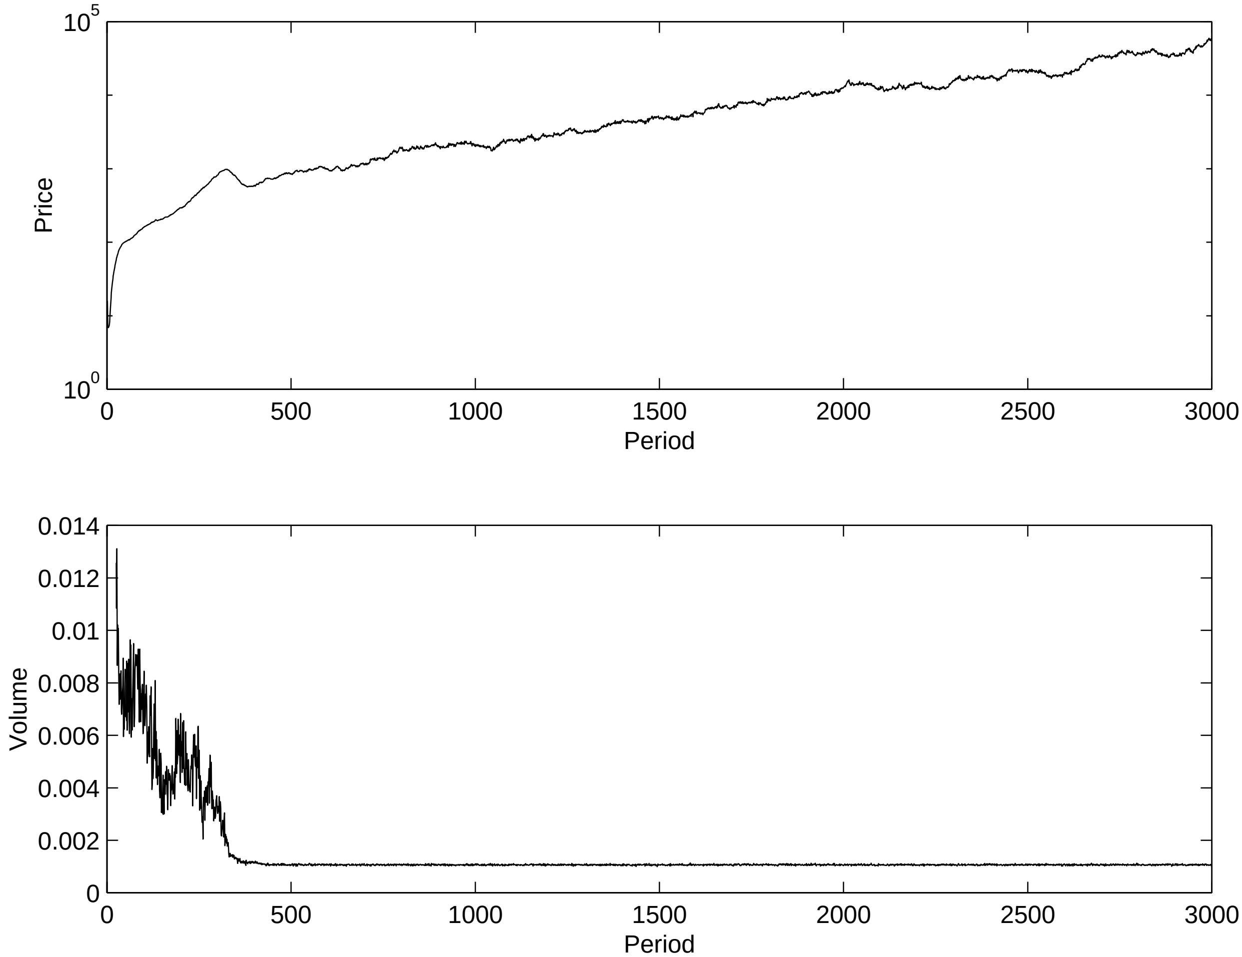

FIGURE 2 Price/Volume: Long Horizons

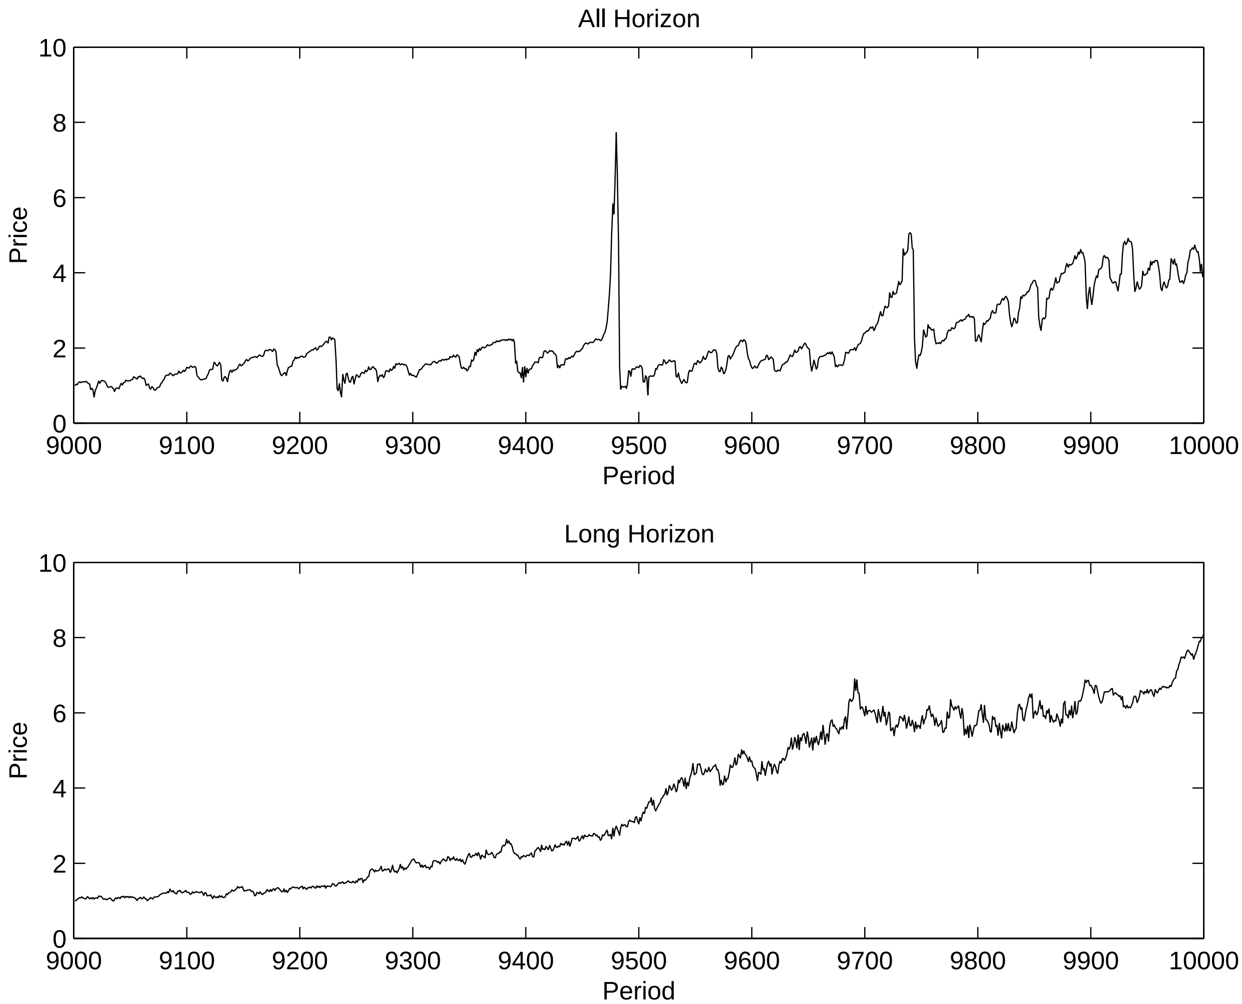

FIGURE 3 Price Comparison

RETURNS, VOLUME, AND VOLATILITY

Figure 4 compares returns across the two different computer experiments. The plot shows monthly returns inclusive of dividends. The return comparisons in Figure 4 confirm the earlier plots showing a much more variable return in the all-horizon case. Visually, returns also appear to exhibit some very large moves, with a few months yielding returns over 50 percent. Also, there appears to be some clumping in that periods of large movements are grouped together. None of these features are present in the lower panel, which corresponds to the long-horizon case. The computer-generated market yields a fairly homogenous set of returns with few, if any, large moves.

Figure 5 presents more evidence on the return distributions. Here, return histograms are compared with normal distributions. In the bottom panel, it is clear that the normal distribution provides a relatively good fit to the computer-generated returns. In the upper panel, the unusual aspects of the large moves in the all-horizon case become clear. The histogram is more peaked and yields several very unusual observations that are well outside the normal range. This replicates the sorts of distributions that are often observed in actual markets.

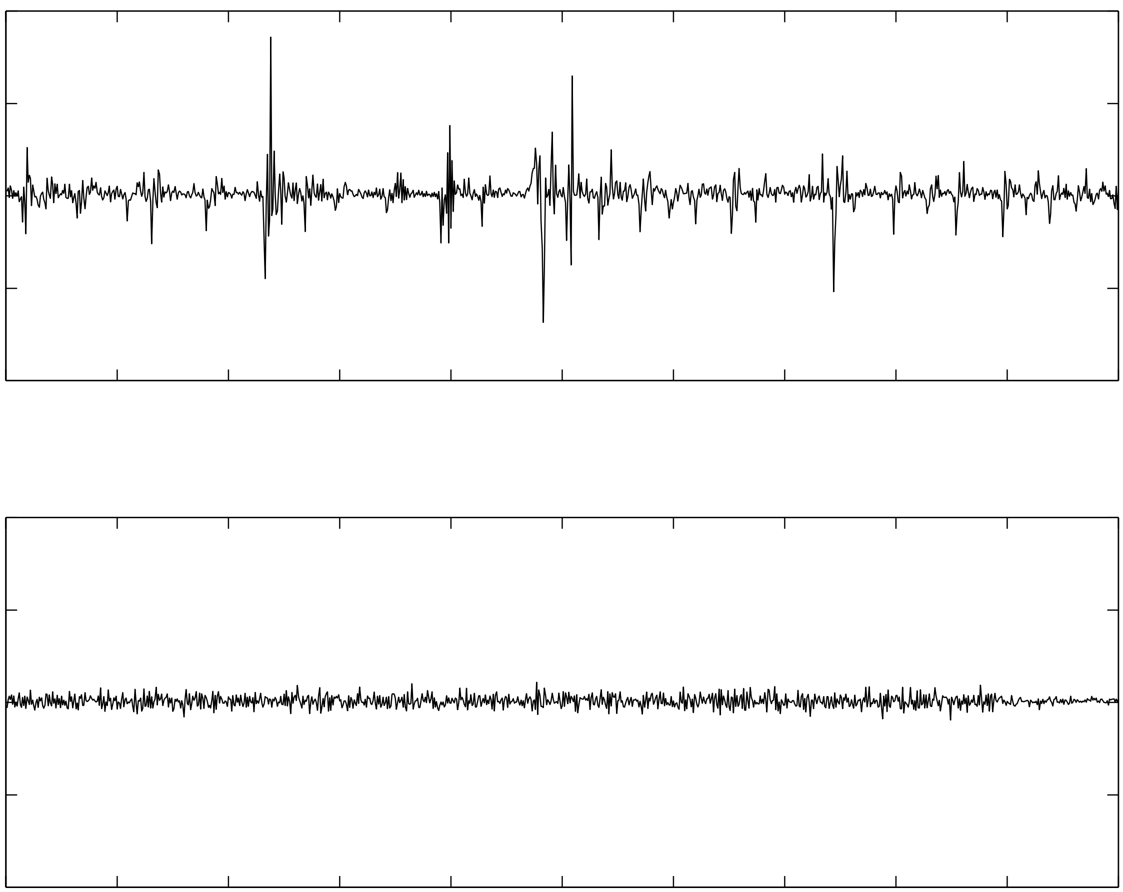

FIGURE 4 Return Comparison

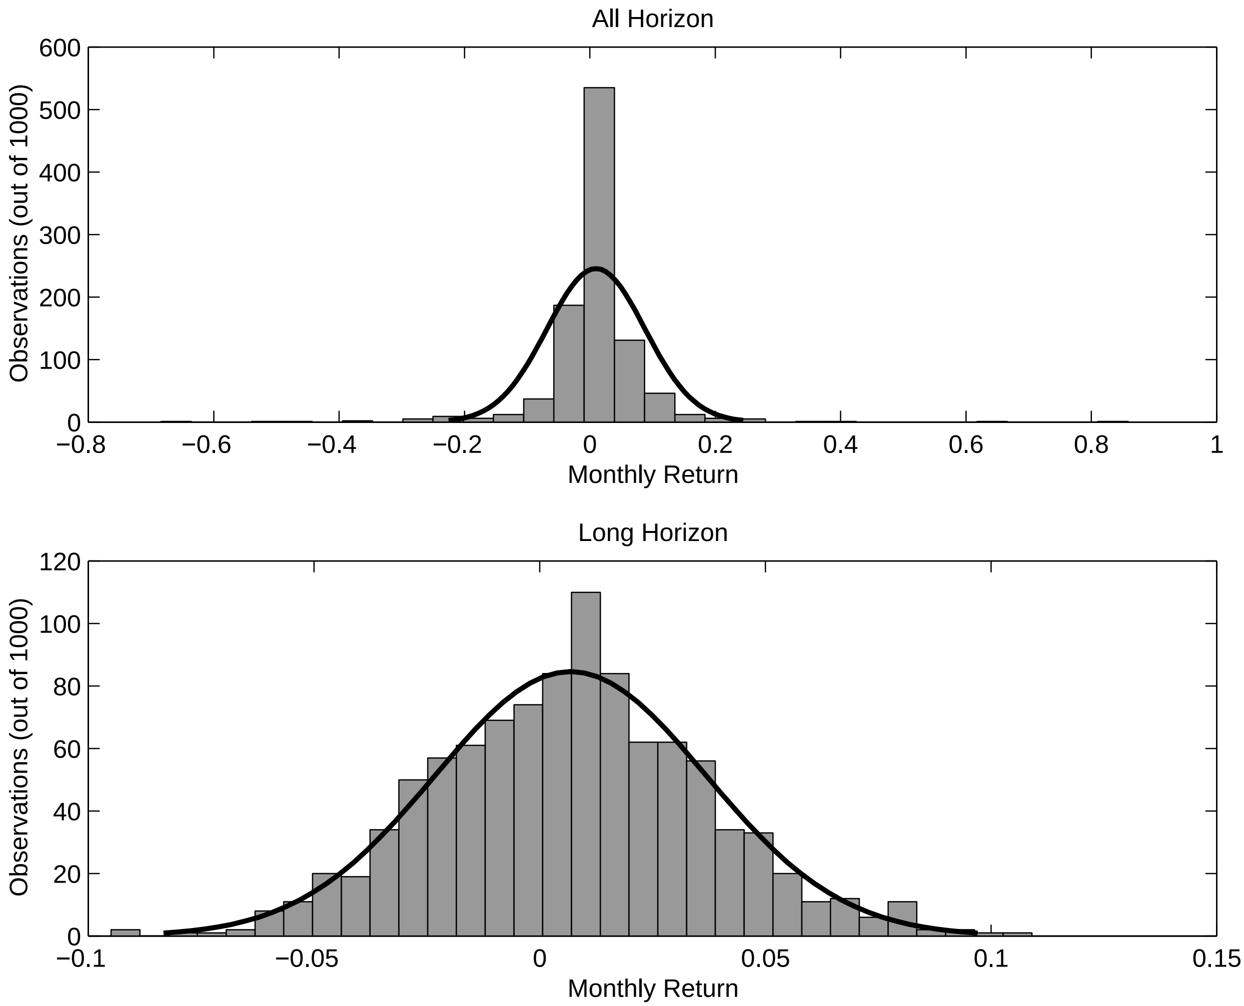

FIGURE 5 Return Distributions

Trading volume for the final 1,000 periods is shown in Figure 6. This plot shows a very strong distinction between the two cases, with very large and fluctuating volume in the all-horizon case and nearly zero volume in the long-horizon case. Obviously, trading volume is an important part of actual financial markets. The lower panel is merely reflecting the fact that the set of agents is in close agreement for asset valuations. In this case, they have no interest in trading with each other, while in the all-horizon case, differences of opinion and the desire to trade do not disappear.

Table 1 summarizes some of the results on the equity returns and trading volume. The mean return is larger in the all-horizon case than the long-horizon case. The standard deviation shows a large increase in moving from the long-horizon to the all-horizon case. Neither case demonstrates any significant skewness in returns. The most dramatic difference appears in the kurtosis estimates. These measure the thick-tailed aspects of the return distribution, which have been displayed in Figure 5. For a normally distributed random variable, kurtosis is 3. For the long-horizon case, it is very close to 3, with an estimate of 3.2. The all-horizon case shows clear evidence for leptokurtosis or "fat tails," with an estimated kurtosis of 32. The volume numbers

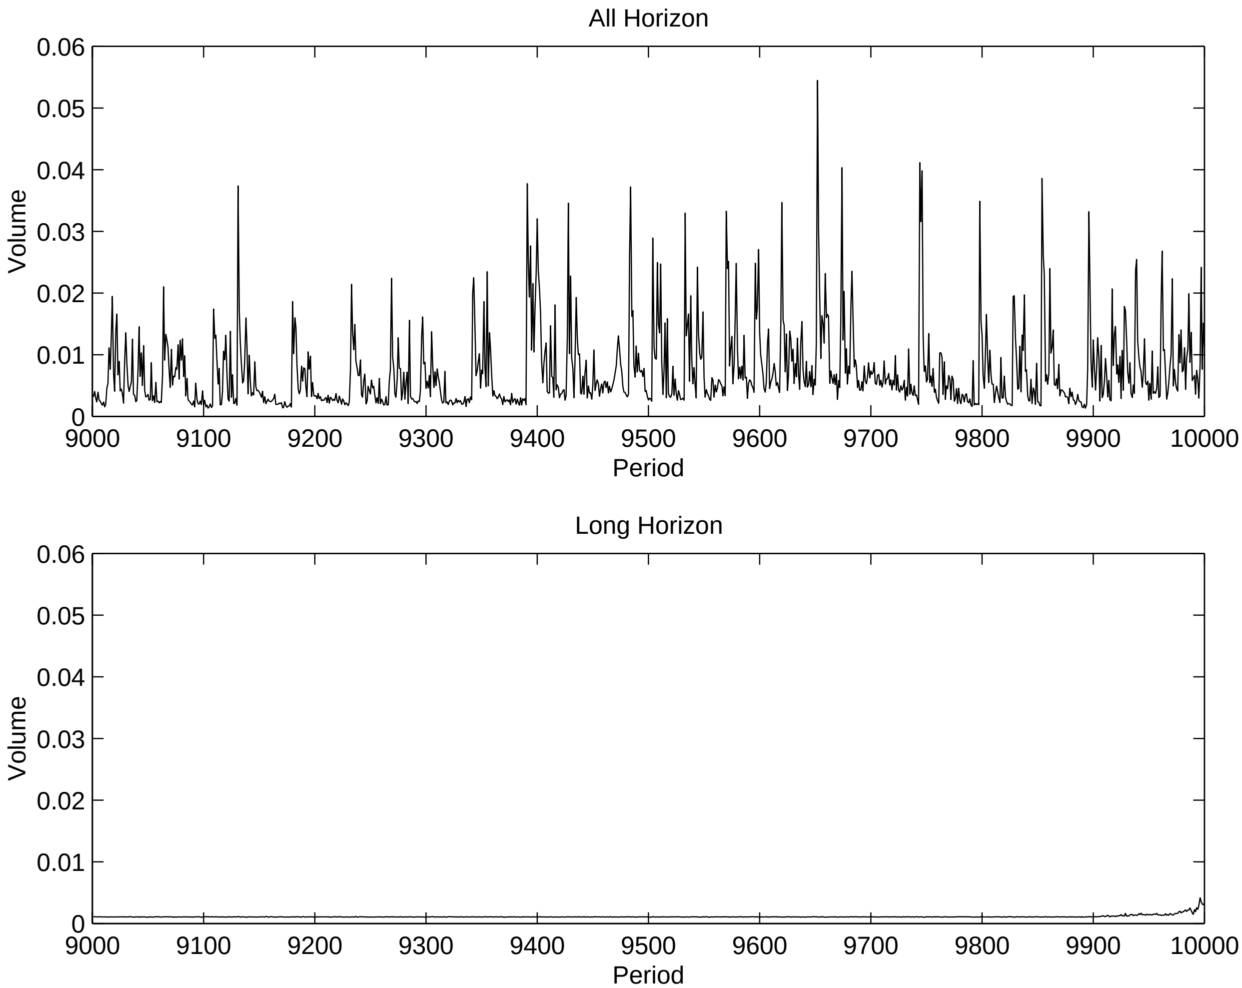

FIGURE 6 Trading Volume

TABLE 1 Excess Return Summary Statistics

	Mean	Std	Skewness	Kurtosis	Volume
All Horizon	10.6	27.1	0.3	32.6	8.7
Long Horizon	6.2	10.4	0.1	3.2	1.3
S&P	8.9	19.6	0.5	12.9	[15,78]

Summary statistics: Mean and Std are the annualized mean and standard deviation of the returns series, inclusive of dividends. Skewness and kurtosis are estimated at the monthly horizon. Values for the S&P are the total return less the 30-day T-bill rate monthly from January 1926 through June 1998. Trading volume reflects the percentage turnover at an annual rate. The value corresponding to the line S&P is the range of NYSE reported values from 1958 through 1999. It is taken from *NYSE Fact Book 1999* (2000).

report turnover at an annual rate. The all-horizon case displays trading volume that is nearly 7 times the value for the long-horizon case. Comparison numbers for the S&P are also presented. On means and standard deviations, the S&P values are between the all-horizon and long-horizon cases. The all-horizon market is actually generating more volatility and more large moves than in actual monthly data. In terms of trading volume, the NYSE clearly generates more turnover per year than the market simulations.[13]

Returns and trading volume in actual markets display several interesting dynamic features that have been hinted at in some of the earlier figures. First, returns are close to uncorrelated. Second, volatility, or the absolute value of returns, is positively correlated. In other words, large moves in either direction tend to follow large moves. Finally, trading volume is also positively correlated. Figure 7 displays the autocorrelations for returns, absolute returns, and trading volume for the long-horizon case. As in real financial data, returns show little or no correlations,

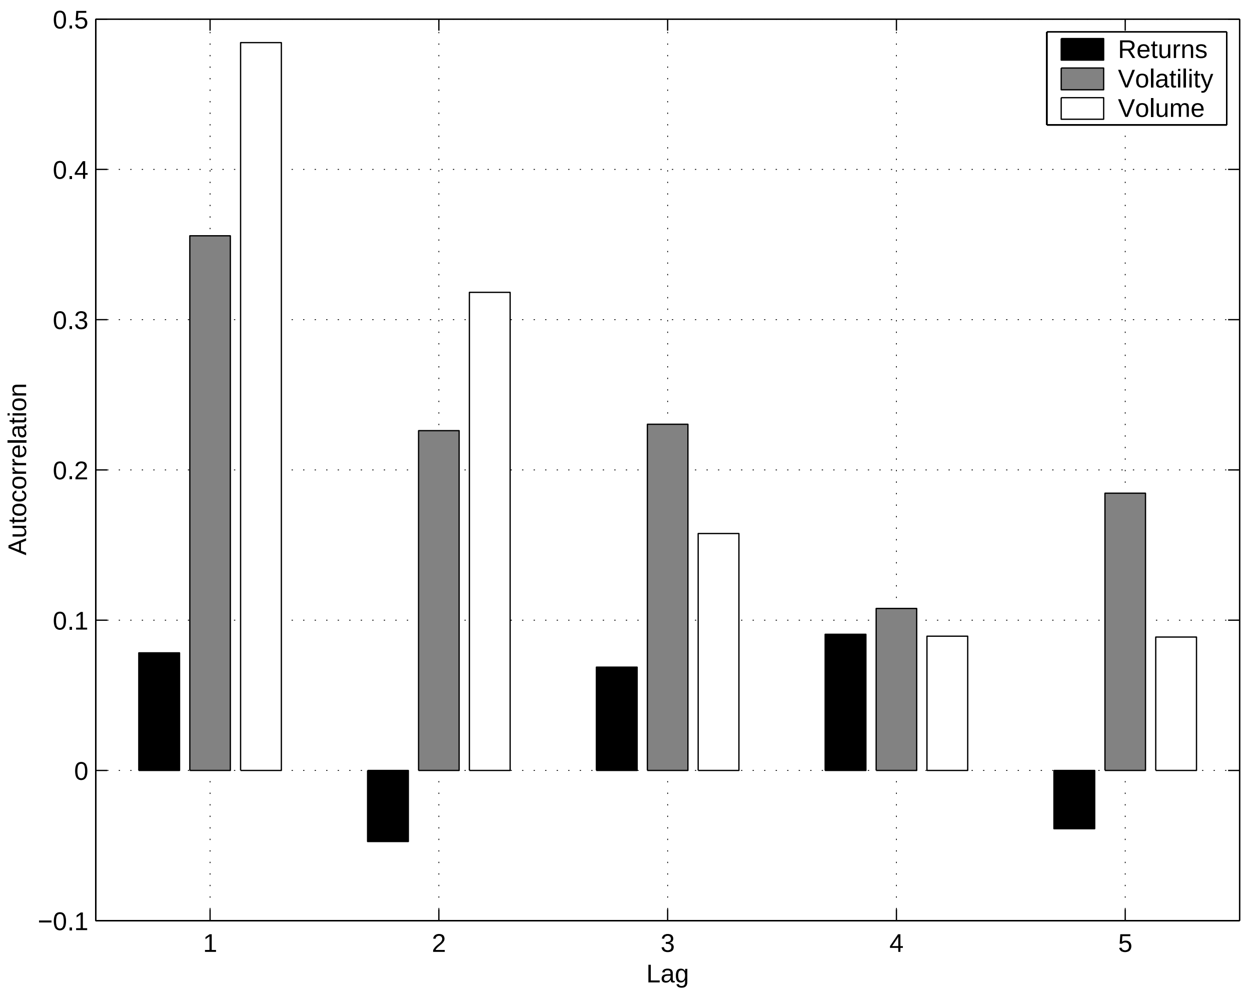

FIGURE 7 All Horizons: Return, Volatility, and Volume Autocorrelations

[13] This number is presented more for information purposes. This two asset market is difficult to compare with the NYSE in terms of trading volume, since the latter obviously has many more opportunities for trade.

positive or negative. However, both volume and volatility display strong positive correlations, going out several periods, as would be the case with actual financial time series. Figure 8 displays the autocorrelations for the market populated with long-horizon investors only. This displays a picture quite different from actual markets. Returns show a small amount of negative correlation at one lag, and then zero after that. The volatility and volume series show only negligible correlation compared with those from the all-horizon experiment. These figures clearly show a very different dynamic in the two different cases, with the all-horizon case showing a picture that more accurately reflects real markets.

Dividend Yields

Another feature of real markets is that they appear to deviate from accepted fundamental models of valuation, yielding some predictability from classical ratios such as price/earnings and dividend yields.[14] In these simple simulated markets, this a very easy and direct experiment as to

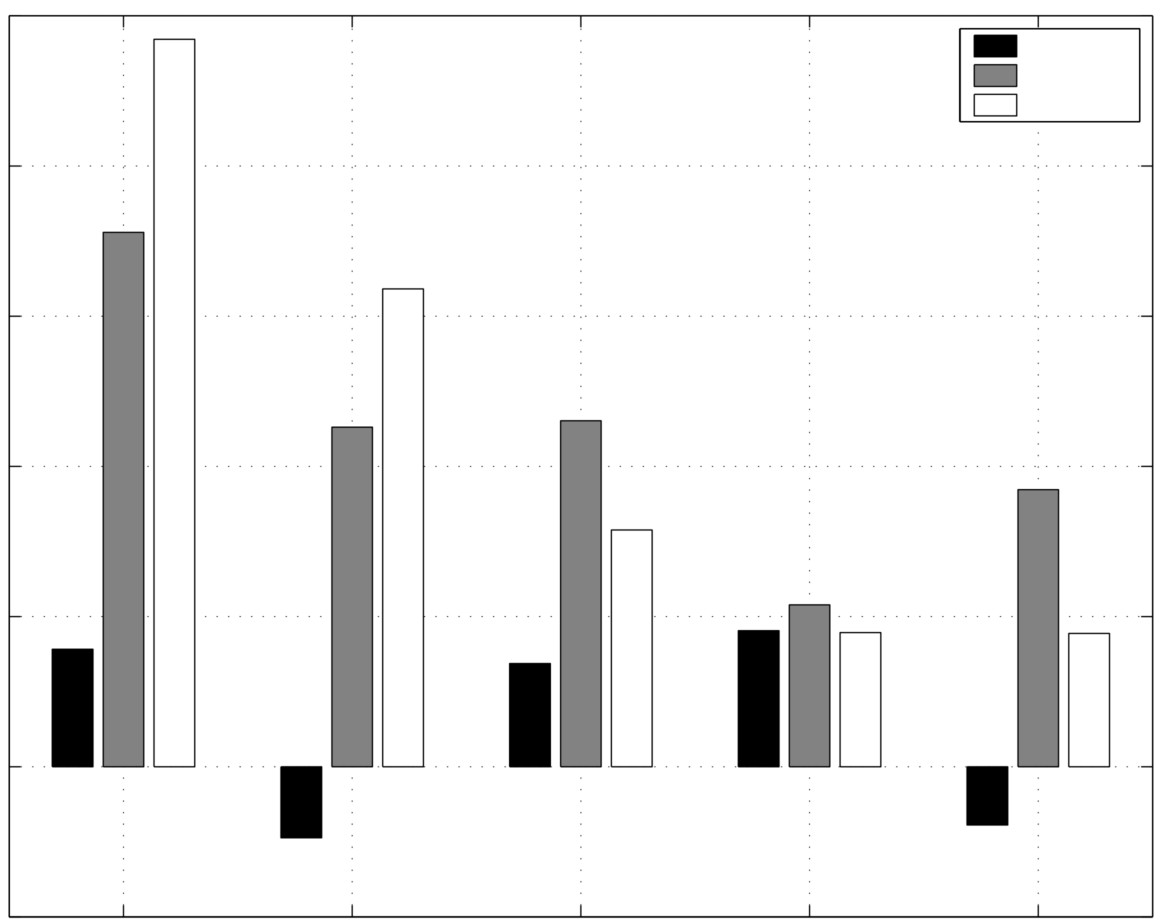

[14] See, for example, Campbell & Shiller (1988), and for a recent commentary on valuation ratios and today's markets, see Campbell & Shiller (2001).

FIGURE 8 Long Horizon: Return, Volatility, and Volume Autocorrelations

how well these markets are doing in terms of reflecting fundamental valuation. Since the equity asset only reflects a stochastic dividend stream, the dividend yield is the ratio of choice for valuation. In a constant growth situation, both the price and dividend will be growing at the same rate, and the ratio should be constant. Figure 9 compares the dividend yields in the two cases. When investors are long horizon in nature, the dividend yield is nearly flat, with only small variation around a value near 5 percent.[15] The situation is very different from the all-horizon case. It is clear that the dividend yield takes some very wide swings, going as high as 12 and as low as 2 percent. It is far from a stable series. This compares much more favorably to dividend yields and price/earnings ratios in actual financial series.

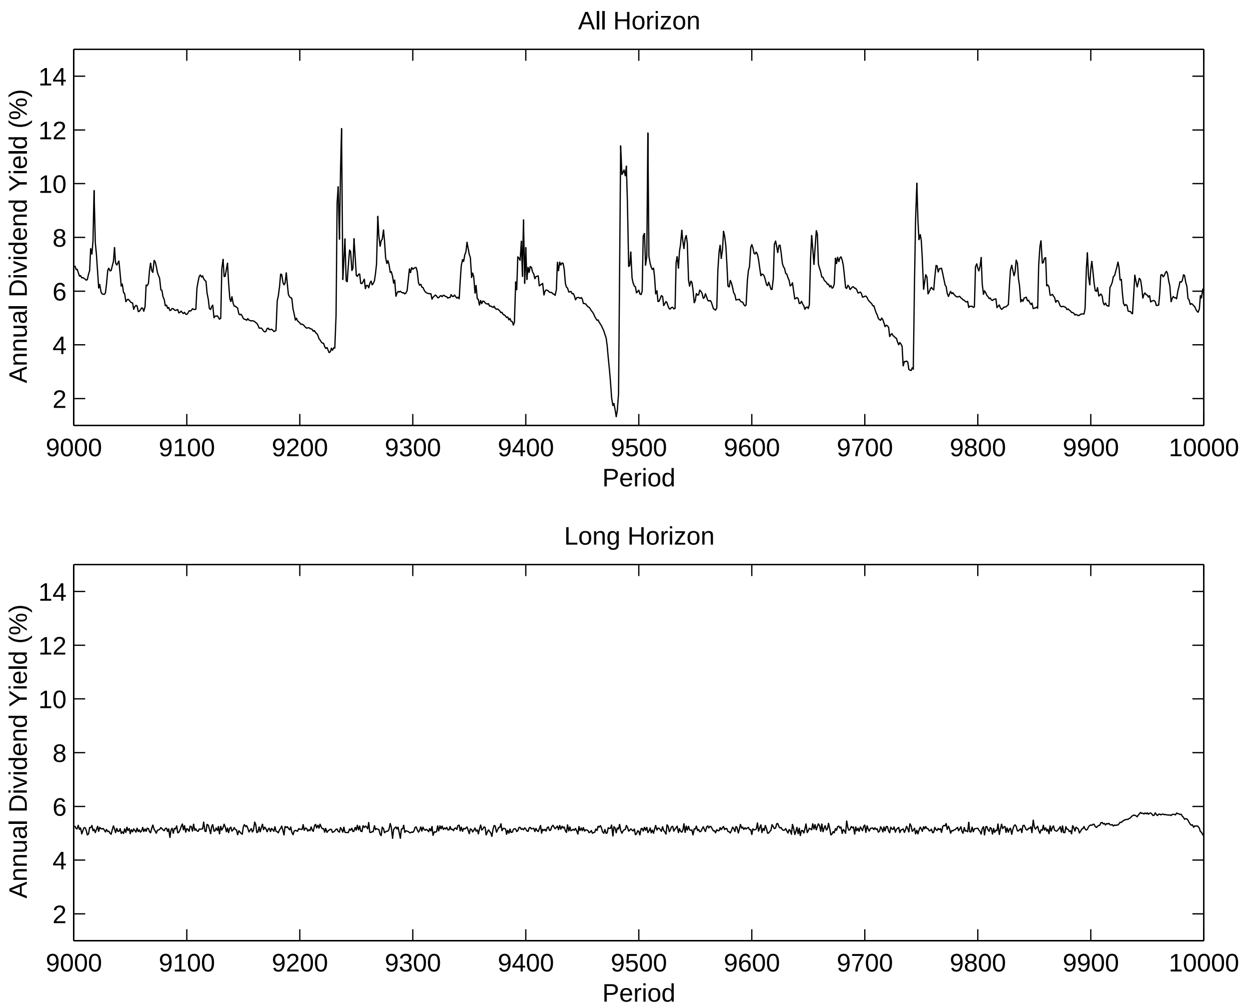

FIGURE 9 Dividend Yields

[15] The monthly dividend yield is annualized by multiplying by 12. The historical average for the S&P is close to 5 percent.

Crash Dynamics

The last two figures present an initial picture of what may be behind the large price changes in this model.[16] They take a snapshot of a short time series of prices, including some large crash periods, and compare these with two other related series.

Two measures will be used that try to assess some aspect of trader heterogeneity. As mentioned earlier, a key aspect of agent-based models is that the actual level of heterogeneity in the market is endogenous. It is possible this may be a precursor to market instability. First, the most obvious magnitude to check around crashes is trading volume. In Figure 10, several crashes are plotted along with the trading volume series. There is a weak indication that trading volume increases greatly after crashes, but it doesn't appear to show a very strong pattern before any of the large price drops in the figure, so it would be difficult to blame crashes on trading volume.

In Figure 11, the same price series is plotted against another measure of agent activity, rule dispersion. This variable measures the fraction of the rules that are currently being used by

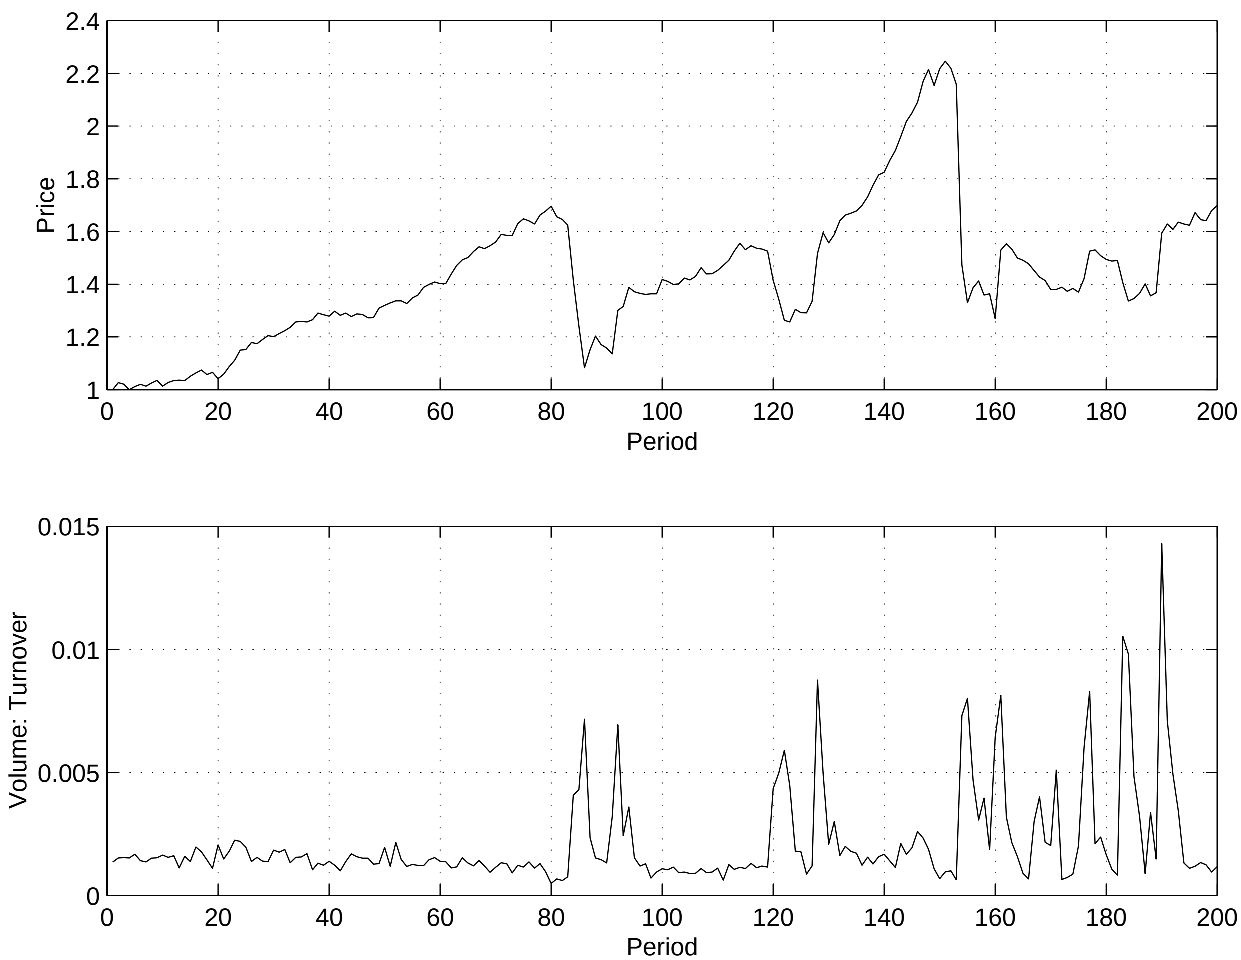

FIGURE 10 Crashes and Volume

[16] For a much more detailed statistical analysis, see LeBaron (2001c).

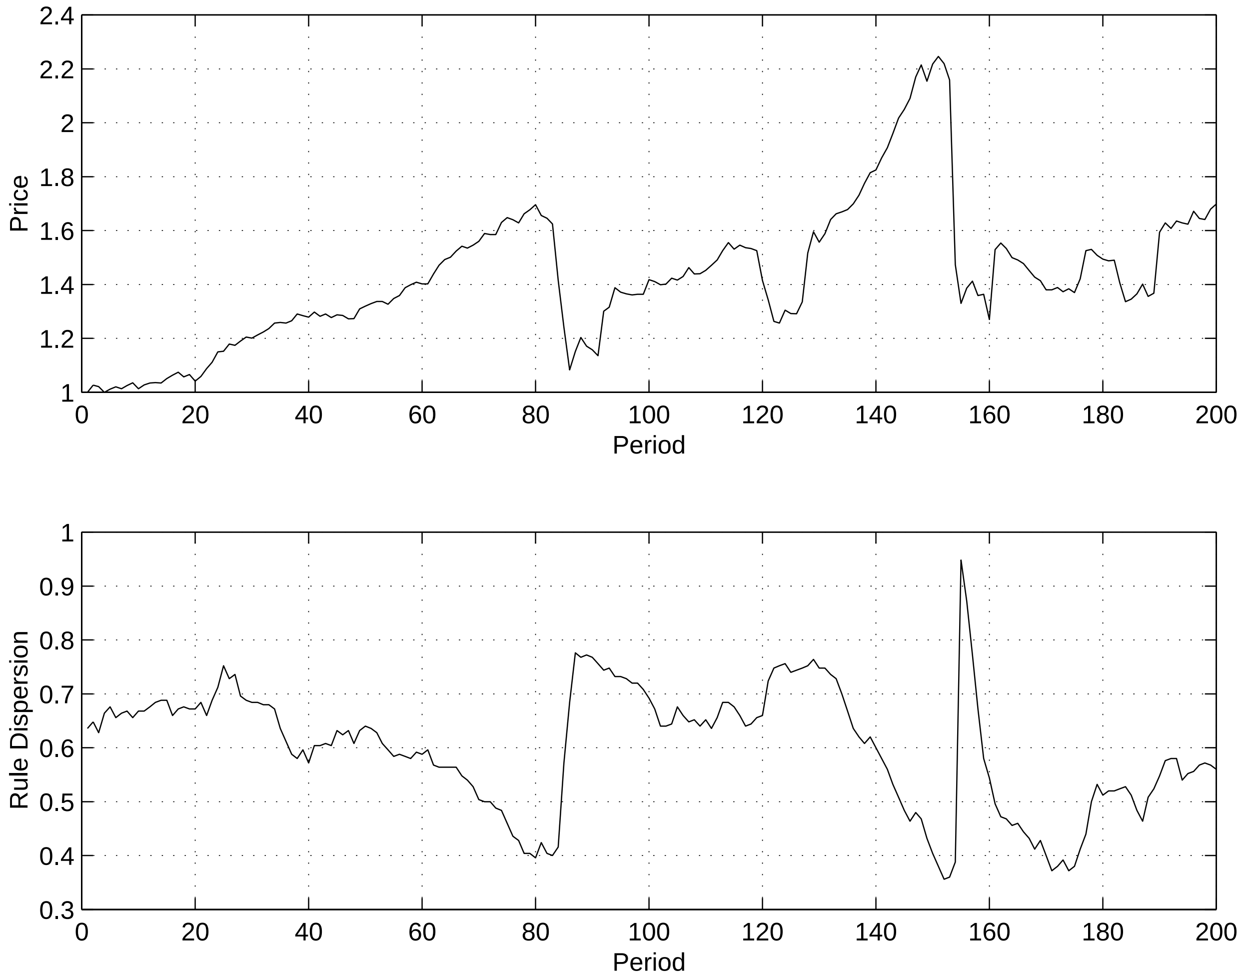

FIGURE 11 Crashes and Dispersion

agents in their active trading strategies. This comparison shows a very dramatic pattern occurring before both large crashes. Rule dispersion begins to fall long before the crash and reaches historically low levels at or near the crash date.

Armed with only this small picture of the overall dynamics, it is difficult to confirm an exact cause for market crashes. However, a simple story is starting to emerge. During the run-up to a crash, population diversity falls. Agents begin to use very similar trading strategies as their common good performance begins to self-reinforce. This makes the population very brittle, in that a small reduction in the demand for shares could have a strong destabilizing impact on the market. The economic mechanism here is clear. Traders have a hard time finding anyone to sell to in a falling market since everyone else is following very similar strategies. In the Walrasian setup used here, this forces the price to drop by a large magnitude to clear the market. The population homogeneity translates into a reduction in market liquidity.[17]

[17] This overall dynamic has some interesting parallels to the problems encountered by Long Term Capital Management. This hedge found it difficult to reduce its positions since many other traders had similar trades in place. See Lowenstein (2000) for a detailed description.

SUMMARY AND CONCLUSIONS

The results in this paper can be summarized along two dimensions. Its most important result is to provide a counter example to the argument that evolutionarily less rational strategies should be driven out of the market. Its second result is in generating time series that appear reasonable for fitting certain difficult-to-replicate features from actual markets.

As a counter example to evolutionary arguments about market efficiency, this model calls into question the basic structure of this argument. Who exactly is “less rational” in a world of heterogeneous agent investors? Will it be clearly obvious to investors to take a long-run perspective in a market dominated by short-run investors? In many ways, these computer experiments may simply be demonstrating that it is very difficult to go against the flow of the current market, even if you feel you must be right. Eventually, your performance relative to others will induce you to take a different (shorter) run perspective, and possibly further add to market instability and deviations from fundamental values.

It is important to realize that the evolutionary experiments implied in Friedman (1953) are quite a bit simpler than the actual market population dynamics. These arguments suggest a well-defined population of super-rational traders that is getting invaded by less rational types. In such a situation, the rational types may have already defined the environment, since their numbers allow them to dominate the pricing of traded assets. In other words, they have defined the rational world. In such a situation, the evolutionary argument is correct, and it should be difficult for the short-run types to stage a successful invasion of the market. The problem is that the more rational types often do not get to start with the luxury of having dominated the market. They must take over in a sea of noise from a heterogeneous population of less-than-rational types. To further complicate matters, this population may be changing its character over time, so the question of optimality gets quite murky.

Several time series features of financial data have proved to be difficult to replicate in standard models of asset pricing. This agent-based simulation easily handles many of these empirical regularities. These include “fat tailed” return distributions, increased volatility and trading volume, volatility and volume persistence, and highly persistent and variable dividend yields. Although the model may not match all of these entirely, it is still impressive that it can attack such a wide range of financial puzzles.

The model also takes a stand on the causes for large moves or market crashes. In the agent-based laboratory, the evidence suggests that crashes are caused by a move toward generally homogeneous markets. As agent strategies become aligned, market liquidity falls, and the market dynamics become brittle as traders cannot find counter parties for their trades. Hopefully, this will lead to further testable hypotheses concerning large price moves and to more general models of market liquidity.

The technology of agent-based financial markets is still in its infancy, but it appears to be a promising route for increasing our understanding of the dynamics of financial markets. It provides new insights for trying to understand financial theories in a world that may be far from the equilibrium that the theories were designed to describe. Features of financial series, such as large moves and excess volatility, may become easier to understand when viewed from a multiagent evolutionary perspective. There is a long way to go in terms of model fitting and assessment, but, for the moment, agent-based theories should take their place as a viable alternative to more traditional financial theories.

REFERENCES

Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. and Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. Durlauf and D. Lane, eds, "The Economy as an Evolving Complex System II," Addison-Wesley, Reading, MA, pp. 15-44.

Blume, L. and Easley, D. (1990), "Evolution and market behavior," *Journal of Economic Theory* **58**, 9-40.

Campbell, J. Y. (1999), Asset prices, consumption, and the business cycle, in J. B. Taylor and M. Woodford, eds, "Handbook of Macroeconomics," North Holland, pp. 1231-1303.

Campbell, J. Y. and Shiller, R. (1988), "The dividend-price ratio and expectations of future dividends and discount factors," *Review of Financial Studies* **1**, 195-227.

Campbell, J. Y. and Shiller, R. J. (2001), Valuation ratios and the long-run stock market outlook: An update, Technical Report 8221, National Bureau of Economic Research.

Cont, R. and Bouchaud, J. P. (2000), "Herd behavior and aggregate fluctuations in financial markets," *Macroeconomic Dynamics* **4**, 170-196.

DeLong, J. B., Shleifer, A., Summers, L. H. and Waldmann, R. (1991), "The survival of noise traders in financial markets," *Journal of Business* **64**, 1-19.

Epstein, J. M. and Axtell, R. (1997), *Growing Artificial Societies*, MIT Press, Cambridge, MA.

Friedman, M. (1953), The case for flexible exchange rates, *in* "Essays in positive economics," University of Chicago Press, Chicago, IL.

Kim, G. and Markowitz, H. (1989), "Investment rules, margin, and market volatility," *Journal of Portfolio Management* **16**(1), 45-52.

LeBaron, B. (2000), "Agent based computational finance: Suggested readings and early research," *Journal of Economic Dynamics and Control* **24**(5-7), 679-702.

LeBaron, B. (2001a), "A builder's guide to agent-based financial markets," *Quantitative Finance* **1**(2), 254261.

LeBaron, B. (forthcoming 2001b), "Empirical regularities from interacting long and short memory investors in an agent based stock market," *IEEE Transactions on Evolutionary Computation*.

LeBaron, B. (2001c), Volatility magnification and persistence in an agent based financial market, Technical report, Graduate School of International Economics and Finance, Brandeis University, Waltham, MA.

Levy, M., Levy, H. and Solomon, S. (1994), “A microscopic model of the stock market: cycles, booms, and crashes,” *Economics Letters* **45**, 103-111.

Lowenstein, R. (2000), *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, Random House, New York, NY.

Lux, T. (1997), “Time variation of second moments from a noise trader/infection model,” *Journal of Economic Dynamics and Control* **22**, 1-38.

Merton, R. (1969), “Lifetime portfolio selection under uncertainty: The continuous time case,” *Review of Economics and Statistics* **51**, 247-257.

NYSE Fact Book 1999 (2000), New York Stock Exchange, New York, NY.

Palmer, R., Arthur, W. B., Holland, J. H., LeBaron, B. and Tayler, P. (1994), “Artificial economic life: A simple model of a stock market,” *Physica D* **75**, 264-274.

Samuelson, P. (1969), “Lifetime portfolio selection by dynamic stochastic programming,” *Review of Economics and Statistics* **51**, 239-246.

Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.

Sargent, T. (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton, NJ.

Simon, H. (1969), *The Sciences of the Artificial*, MIT Press, Cambridge, MA.

Tay, N. S. P. and Linn, S. C. (2001), “Fuzzy inductive reasoning, expectation formation and the behavior of security prices,” *Journal of Economic Dynamics and Control* **25**, 321-362.

PRICE EFFICIENCY AND RISK SHARING IN TWO INTER-DEALER MARKETS: AN AGENT-BASED APPROACH

J. YANG, Bank of Canada*

ABSTRACT**

This paper investigates the effect of institutional arrangements on market performance using agent-based market simulation. The direct inter-dealer market is modeled as a decentralized dealership market and the inter-dealer broker market, is modeled as an order book market with double auction algorithm. On the simulated markets, six dealers trade with customers first and then with each other on a direct inter-dealer market and an electronic broker market, respectively, in two experiments. Market quality is investigated based on the results of the simulations. We find that given the same level of post-trade transparency, price discovery is relatively faster in an electronic broker market than in a direct dealer market, while the direct dealer market provides for greater opportunity for risk sharing. Furthermore, post-trade transparency increases price efficiency in both inter-dealer markets.

INTRODUCTION

This paper addresses the issue of institutional design in an inter-dealer market. We focus on inter-dealer trading and compare the effects of two distinct inter-dealer markets on market performance. These two markets are the direct inter-dealer market and the brokered inter-dealer market.

We directly compare the effects of two market structures on price efficiency and inventory control. The two markets differ only in the way trading is organized. In a direct inter-dealer market, trades are bilateral, simultaneous, and decentralized. Trades can occur at different prices at the same time. In an electronic broker market, trades are continuous and centralized and can only occur at one price at the same point in time, which is either at the best available bid or ask.

The experimental approach to financial markets is a growing field. A few experimental studies explore market performance in the context of the dealership market. Flood et al. (1999) examine the effects of pretrade transparency on price discovery. Their key comparison is between fully public price queuing (a pretrade transparent market) and bilateral quoting (a pretrade opaque market). They find that market liquidity is lower in the opaque market, but price discovery is faster in the pretrade opaque market. Thus, a trade-off exists between price

* *Corresponding author address*: Jing Yang, Bank of Canada, Financial Markets Department, 234 Wellington St., Ottawa, K1A 0G9; e-mail: jyang@bank-banque-canada.ca.

** This is a shortened version of Yang (2001). The author is grateful to all the participants of the Agent 2000 workshop at The University of Chicago and CEF2000 in Barcelona for their comments. Any errors that remain are the author's responsibility.

efficiency and liquidity. Bloomfield and O'Hara (1999) investigate the effects of both pretrade transparency and post-trade transparency on market performance. They find that post-trade transparency increases price efficiency but also increases the bid/ask spread. They find that pretrade transparency has no discernible effects on market performance. Sharing similar interests with these two studies, we attempt to use an agent-based approach in this paper, rather than a human-subject-based experiment, to study questions related to market institutional design.

THE MODEL AND THE EXPERIMENTS

Consider an economy with two assets, one riskless and one with stochastic liquidation value ("true value"), representing foreign exchange (FX). Dealers trade with customers first, and then they go to one of the inter-dealer markets, either a decentralized direct inter-dealer market or a centralized inter-dealer broker market to unwind their undesired inventory. The risky asset trades in periods $t = 1, 2, \ldots,$ T. Each trading period contains a series of trading rounds, *round* = *1, 2,...,* R. The full information price of FX at time t is denoted by F_t, which is composed of a series of increments (e.g., interest differentials) so that

$$F = \sum_{i=0}^{T} r_i \ .$$

The increments r_i are i.i.d. with mean zero. Each increment r_i is realized immediately after each trading in period t. Realizations of the increments represent the flow of public information over time. The value of FX at t is defined as

$$F_t = \sum_{i=0}^{t} r_i \ .$$

Three signals define each period's information environment prior to dealer i's quote. The first two signals are received simultaneously. Before the third signal arrives, the first two signals are:

$$\begin{aligned} &S_t = F_t + \varepsilon_t, \text{ and} \\ &C_{i,}t = F_t + \omega_{i,t}. \end{aligned} \tag{1}$$

The noise terms ε_t and ω_t are normally distributed about mean zero, are independent of one another and across periods, and have variance σ_ε^2 and σ_ω^2, respectively. At the outset of each period t, all dealers receive a public signal S_t about the full-information value F, and also a private signal $C_{i,t}$. One potential source of private signals at the FX dealer level is order flow from non-dealer customers. Each dealer has sole knowledge about his own customers' order flow. To the extent that this flow conveys useful information, it potentially can be exploited in inter-dealer trading.

At the end of each trading round, dealers observe market information that represents a signal of market-wide order flow. This is given by:

$$V_t = \sum_{i=1}^{n} T_{i,t} + \zeta_t \ . \tag{2}$$

This net order flow measures the difference in buy and sell orders, and $T_{i,t}$ is the order placed by dealer i. The noise term ζ_t is normally distributed with mean zero and variance σ_ζ^2, which represents the precision of the signal. The empirical analogue of V_t is the signed order-flow information communicated by broker intercoms. The signal V_t plays a very important role in this paper. We interpret the precision of this order flow, σ_ζ^2, as the degree of transparency in the market. In later experiments, we vary the value of the precision in this signal to vary the level of transparency on FX markets.

In this paper, we model two inter-dealer markets. The direct inter-dealer market is a decentralized, quote-driven market. The electronic broker market is a centralized, order-driven, and continuous double auction market. Dealers trade with customers, and then they go to the inter-dealer market and trade among themselves. The choice of inter-dealer market is exogenous.[1]

For the direct inter-dealer market, dealers have to call one another bilaterally to ask for a quote. Then dealer *j* lays off his undesired inventory at dealer *i*'s price. In this market, quoting is simultaneous and independent, which implies that dealer *i*'s quote is not conditioned on the other dealers' quotes. Dealer *i* can only observe dealer *j*'s quote, not any other's quote.[2] This means that quote search is costly. Quotes are good for any size order. Dealers cannot refuse to quote, which implies that each dealer has an obligation to make a market. No trade information is revealed during the trading round. Only at the end of each trading round can each dealer observe a noisy signal about market-wide net order flow.

For the inter-dealer broker market, we choose a double auction setting that is closely related to Yang (2000).[3] In our simplified double auction market, dealers can either submit a bid/ask, or accept an existing bid/ask. A transaction occurs when an existing bid/ask is accepted. At the beginning of each trading round, we suppose a random permutation of the dealers, which determines the subsequent order of dealers. Initially, the dealers come to the market with their own price expectation, and they attempt to post or accept a bid (ask) order by comparing their price with the existing best ask (bid). They can only observe the best bid and the best ask price.

---

1 An interesting future research topic is to make the choice between two inter-dealer markets endogenous to dealers. In reality, a common feature of dealership markets is that dealers have a choice between direct dealer trading or going to a broker market. Dealers are observed to switch between the markets, but the factors that determine such switching would be very interesting to investigate.

2 To keep relatively similar information sets on both inter-dealer markets for comparability, we restrict the visibility of all the quotes in the direct inter-dealer market. In future research, this rule can be relaxed and each dealer can observe a few other dealers' quotes and trade with the lowest one, or not trade at all but just extract information.

3 A few extensions in auction setting from Yang (2000) are first, that the assumption of fixed order size is relaxed, so dealers can submit an order at any size, and second, that partial execution of orders becomes possible.

Let a be the best ask, b the best bid, $\hat{u}_{i,t}$ the expectation of the true price in this period, and γ the bid-ask spread. Then dealers trade with each other according to the following scenarios.

Scenario 1. *If a best bid,* b*, and a best ask,* a*, exist in the market*

- If $\hat{u}_{i,t} > a$, he will submit a market order, buy $\delta(\hat{u}_{i,t} - a)$ units at this ask price;[4]

- If $\hat{u}_{i,t} < b$, he will submit a market order, sell $\delta(b - \hat{u}_{i,t})$ units at this bid price;

- If $b < \hat{u}_{i,t} < a$ and $\hat{u}_{i,t} < (a + b)/2$, he will post a buy order at a price of $(\hat{u}_{i,t} + \gamma)$ and a size of $\delta\gamma$.

- If $b < \hat{u}_{i,t} < a$ and $\hat{u}_{i,t} > (a + b)/2$, he will post a buy order at a price of $(\hat{u}_{i,t} - \gamma)$ and a size of $\delta\gamma$.

Scenario 2. If only the best ask, a*, exists*

- If $\hat{u}_{i,t} > a$, he will submit a market order, buy $\delta(\hat{u}_{i,t} - a)$ at this ask price;

- If $\hat{u}_{i,t} < a$, he will post a buy order at a price of $(\hat{u}_{i,t} - \gamma)$ with a size of $\delta\gamma$.

Scenario 3. If only the best bid, b*, exists*

- If $\hat{u}_{i,t} < b$, he will submit a market order, sell $\delta(b - \hat{u}_{i,t})$ at this bid price;

- If $\hat{u}_{i,t} > a$, he will post a sell order at a price of $(\hat{u}_{i,t} + \gamma)$ and a size of $\delta\gamma$.

Scenario 4. If no bid and ask exist,

- He will have an equal chance to post a buy or a sell order at price of $(\hat{u}_{i,t} - \gamma)$ or $(\hat{u}_{i,t} + \gamma)$, respectively, and a size of $\delta\gamma$.

Dealers use a prototypical inventory model to decide a quote, where price is linearly related to the dealer's current inventory,

---

[4] The order size is decided by equation (5), given the current market price.

$$P_{i,t} = \hat{u}_{i,t} - \alpha(I_{i,t} - I^*_{i,t}) + \gamma D_t \tag{3}$$

where $\hat{u}_{i,t}$ is the expectation of F_t conditional on the information available to dealer i at t; $I_{i,t}$ is dealer i's current inventory position; and $I^*_{i,t}$ is dealer i's desired inventory position. The inventory-control effect, α, is generally a function of the relative interest rate and the cost to carry a position, which remains constant in our experiments. The term $(I_{i,t} - I^*_{i,t})$ captures the inventory imbalance. The higher is a positive (negative) inventory imbalance, the lower (higher) is the ask (bid) price a dealer is willing to post. Dealers use quotes as an important tool to perform inventory control. D_t is a direction indicator variable that equals 1 when the transaction price $P_{i,t}$ is the ask and −1 when $P_{i,t}$ is a bid. For a given expectation $\hat{u}_{i,t}$, γD_t picks up the bid-ask half-spread.

Dealer i's quote schedule is a function of the expectation of F_t at the time of quoting, which is denoted $u_{i,t}$. In turn, this expectation is conditioned on the signals S_t and $C_{i,t}$, which dealers observe at the beginning of the period, and which are updated subsequently with newly revealed market information V_t. Each dealer uses a *neural network* to find a good mapping from the three signals. Dealers update their expectations during each trading round. These expectations play a central role in determining dealer i's quote, which we describe next.

We simulate three different transparency settings for each of the two inter-dealer markets by varying the precision of post-trade order flow information. For the semi-opaque setting, the semi-transparent setting, and the perfectly transparent setting, the variance in order flow signal equals 4, 1, and 0, respectively.

RESULTS

In this section, we test three hypotheses concerning the effect of different market structures on price discovery and inventory control. The test of hypothesis 1 is that market structure has significant effect on the price discovery process. Following Bloomfield and O'Hara (1999), we use the change in the price error to measure price efficiency, which is calculated as the absolute value of the difference between the market prices and the theoretical prices. Table 1 gives a descriptive summary for the movement of price errors in the two inter-dealer markets. The results show price errors decline more rapidly in the electronic broker market. In Table 2, the results of a t-test on two sample mean of the convergence rate are reported. The first two P-values suggest to reject the null at the 5% and 10% significance level, respectively, which implies that prices on a dealership market and inter-dealer broker market converge at a different rate.

This suggests that electronic broker markets reveal information more rapidly and completely than direct dealer markets. Since the two inter-dealer markets only differ in the trading mechanism, and electronic broker market trades always happen at the best bid or ask price, we infer that the difference in price efficiency may benefit from the centralized price information, which reduces the searching cost.

The test of hypothesis 2 is that post-trade transparency increases price efficiency. To test this hypothesis, we vary the transparency level on both markets. Since the market-wide net order flow is the only public signal revealed after each trading round, we vary the noise term in this order flow signal to accommodate different post-trade transparency levels. For semi-opaque, semi-transparent and transparent markets, the variance of the noise in the order flow signals equals 4, 1, and 0 respectively. We present the results for the different transparency levels in Table 2. The values of decline from 0.64 under the semi-opaque setting to 0.61 under the

TABLE 1 Change in the Price Errors on Two Inter-Dealer Markets

Price Errors	Average of All Trading Rounds	Average of First 40 Trading Rounds	Average of Last 40 Trading Rounds
E-broker market	0.3838	0.7385	0.2588
	(0.1170)	(0.2774)	(0.0290)
Direct inter-dealer market	0.6856	1.0662	0.568
	(0.1321)	(0.1666)	(0.0151)

Notes: Every experiment is run for 10 periods and 20 trading rounds in each period. Overall, the same experiment is run for 25 times. The statistics shown in this table are the average over 25 runs. Numbers in parentheses are standard errors estimated using the 25 runs.

TABLE 2 Impact of Market Structure on Price Efficiency at Different Transparency Levels

Price Efficiency Measure	β_1	β_2	**H0: $\beta_1 = \beta_2$**
Semi-opaque setting	0.6412	0.7810	-2.5792
	(0.182)	(0.201)	[**0.006**]
Semi-transparent setting	0.6120	0.6531	-1.4837
	(0.081)	(0.012)	[**0.072**]
Perfect post-trade transparency	0.4206	0.4418	-0.4028
	(0.160)	(0.1812)	[0.334]

Notes: The results presented in this table are for the regressions from equation $Y_{i,t} = \mu_i + \beta_i Y_{i,t-1} + \nu_{i,t}$. The dependent variables are the price errors of the direct inter-dealer market or broker market. Numbers in parentheses are the standard errors of the estimates using the 25 runs. The last column is the statistic for a t-test, with the P-values reported in the brackets. Every experiment is run for 10 periods and 20 trading rounds in each period. The same experiment is repeated 25 times. The statistics shown in this table are averaged over the 25 runs.

semi-transparent setting, and then to 0.42 under the perfect transparency setting on the electronic broker market. A similar pattern appears in the direct inter-dealer market. This suggests post-trade transparency increases the price efficiency and accelerates the decline of price errors in both markets. This is evidence in support of our hypothesis 2. This result is consistent with the experimental findings of Bloomfield and O'Hara (1999), where they find trade disclosure increases the price discovery process.

The test of hypothesis 3 is that the direct dealer market provides a better inventory control opportunity. To investigate which inter-dealer market offers a better inventory control opportunity, we first examine whether the inventories exhibit mean reversion and then measure the adjustment speed of inventory imbalance corrections. A higher adjustment speed implies better risk sharing and inventory control opportunities for an inter-dealer market.

In inventory models of dealership markets,[5] the degree of competitiveness of a dealer's quotes depends on his relative inventory position as reflected in equation (4). The higher the positive (negative) inventory imbalance, the lower (higher) is the ask price that the dealer is willing to post. When a dealer has an extreme inventory position, he is able to post competitive quotes on one side of the market, and he stands a better chance of executing the public order flow in the desired direction. This results in a relatively quick reduction of his inventory imbalance. On the other hand, when a dealer's inventory is closer to the desired level, he is not able to post competitive prices and therefore stands a poor chance of executing the public order flow. As a result, his inventory takes a longer time to revert to the desired level. This implies that in competitive dealership markets, the relative inventories of the dealers should be mean reverting.

Let $IM_{i,t} = I_{i,t} - I^*_{i,t}$ denote the inventory position of market maker i relative to the desired inventory level at time t. We run the following regression,

$$\Delta IM_{i,t} = a + \theta IM_{i,t-1} + \tau_{i,t}, \tag{5}$$

where i denotes the dealer, and $IM_{i,t}$ represents the inventory imbalance of dealer i at the end of time t. The coefficient θ captures the intensity of mean reversion.

The results, which are reported in the second and third columns of Table 3, show that dealer's inventories in the direct dealer market have greater force of mean reversion (the absolute value of the coefficient). The average mean reversion coefficient for dealers' inventories in the electronic broker market is -0.064, and in the direct dealer market, -0.28.

The speed of adjustment of inventories is directly related to the mean reversion coefficient θ, which represents the fraction of the deviation between actual and desired inventories that is eliminated each day. A useful measure of adjustment speed is the inventory half-life, denoted by h. It is defined as the expected number of days required to reduce a desired inventory by 50%, where

[5] See Ho and Stoll (1983), Madhavan and Smidt (1993), and Hansch, Naik, and Viswanathan (1998) for a theoretical inventory model and empirical tests.

$$h = \frac{\ln 2}{\ln(1-\theta).} \tag{6}$$

The last two columns in Table 3 provide the estimates of the inventory half-life for six dealers in the two inter-dealer markets. The results show that it takes, on average, 12.94 trading rounds for an inventory imbalance to be reduced by 50% in the direct dealer market. In contrast, the inventory half-life, on average, is 22.09 trading rounds in the electronic broker market. This implies, on average, that dealers in the direct inter-dealer market are able to adjust their inventory imbalance faster, which, in turn, implies that dealers carry lower inventory-risk in this market.

CONCLUSION AND FUTURE RESEARCH

In this paper, we directly compare two different types of inter-dealer markets. Given a fixed level of market transparency, we find that there is a trade-off between price efficiency and inventory control. The centralized electronic broker market has greater price efficiency than a decentralized direct dealer market, but the direct dealer market provides greater market depth and a shorter inventory half-life. Post-trade transparency increases the price efficiency of both markets. We close with a discussion on further research. Our experiments investigate the effects of market structure on market performance. Having established substantial differences in market behavior, the natural question follows whether one type of inter-dealer market is likely to dominate another. When the choice of trading on certain inter-dealer markets become endogenous, which inter-dealer market prevails? Such a topic is well suited to exploration in an experimental market setting. We believe further research on this topic will provide more insights into inter-dealer markets.

TABLE 3 Inventory Mean Reversion Coefficients and Implied Half-Life

Dealers	θ on e-Broker Market	θ on Direct Dealer Market	Half-life on e-Broker Market	Half-life on Direct Dealer Market
Dealer 1	-0.0184	-0.0118	38.076	59.091
Dealer 2	-0.0134	-0.1085	52.105	6.7281
Dealer 3	-0.0595	-0.2304	11.989	3.3429
Dealer 4	-0.0460	-0.3458	15.410	2.3333
Dealer 5	-0.0646	-0.8960	11.070	0.0903
Dealer 6	-0.1851	-0.1216	4.0812	6.0388

Note: The results presented in this table are for the regressions based on equation (5) and from the calculation using equation (6). In both markets, the transparency parameter takes the value of 1, which implies that the results shown in this table are for semi-transparent markets. Every experiment is run for 10 periods and 20 trading rounds in each period. Overall, the same experiment is run 25 times. The statistics shown in this table are the average over 25 runs.

REFERENCES

Bloomfield, R., and O'Hara, M. (1999) Market transparency: Who wins and who loses? *The Review of Financial Studie*s, Spring 1999 Vol. 12 pp. 5-35.

Flood, M., Huisman, R., and Mahieu, R. (1999) Quote disclosure and price discovery in multiple-dealer financial markets, *The Review of Financial Studie*s, Spring 1999 Vol. 12 pp. 37-59.

Flood, M., (1994) Market structure and inefficiency in the foreign exchange market, *Journal of International Money and Finance* (1994), 13(2), pp. 131-158.

Garcia, R., and Gencay, R. (1999) Pricing and hedging derivative securities with neural networks and homogeneity hint, *Journal of Econometric*s, forthcoming.

Glosten, L. (1994) Is the electronic open limit order book inevitable? *The Journal of Financ*e, Vol. XLIX, No. 4, September 1994.

Hansch, O., Naik, N., and Viswanathan, S. (1998) Do inventories matter in dealership markets? Evidence from the London Stock Exchange, *The Journal of Financ*e, Vol. LIII, No. 5, 1998.

Hertz, J., Krogh, A., and Palmer, R. (1991) *An Introduction to the Theory of Neural Computation*, Addison-Wesley Publishing Company, Reading, Mass.

Ho, T., and Stoll, H. (1983) The dynamics of dealer markets under competition, *Journal of Finance* 38 (1983) pp. 1053-1074.

Hutchinson, J., Lo, A., and Poggio, T. (1994) A nonparametric approach to pricing and hedging derivative securities via learning networks, *Journal of Finance* 49(3), pp. 851-889.

LeBaron, B., Arthur & Palmer (1999) Time series properties of an artificial stock market, *Journal of Economic Dynamics and Control,* 23 (1999) pp. 1487-1516.

LeBaron, B. (forthcoming 1999) Agent based computational finance: Suggested reading and early research, *Journal of Economic Dynamics and Control.*

Lyons, R. (1997) A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* Vol. 42 (1997) pp. 275-298.

Lyons, R. (1995) Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics*, Vol. 39 (1995) pp. 321-351.

Lamoureux, C., and Schnitzlein, C. (1997) When it's not the only game in town: The effect of bilateral search on the quality of a dealer market, *Journal of Finance*, Vol. LII, No. 2, June 1997.

Madhavan, A. (1992) Trading mechanisms in securities markets, *Journal of Finance*, Vol. XLVII, No. 2, June 1992.

Madhavan, A., and Smidt, S. (1993) An analysis of changes in specialist inventories and quotations, *Journal of Finance*, Vol. XLVIII, No. 5, December 1993.

O'Hara (1995) *Market Microstructure Theory*, Blackwell Publishers, Cambridge, USA.

Yang, J. (2000) The efficiency of an artificial double auction stock market with neural learning agents, in Chen, S.H. (ed.), *Evolutionary Computation in Economics and Finance*, Springer-Verlag (forthcoming).

Yang, J. (2001) Price Efficiency and Inventory Control in Inter-dealer Trading, Bank of Canada, internal report.